arXiv:1007.1328v1 [math.CO] 8 Jul 2010

On Belief Propagation Guided Decimation for Random k-SAT

Amin Coja-Oghlan*
University of Warwick, Mathematics and Computer Science,
Zeeman building, Coventry CV4 7AL, UK
`a.coja-oghlan@warwick.ac.uk`

July 9, 2010

Abstract

Let $\mathbf{\Phi}$ be a uniformly distributed random k-SAT formula with n variables and m clauses. Non-constructive arguments show that $\mathbf{\Phi}$ is satisfiable for clause/variable ratios $m/n \leq r_k \sim 2^k \ln 2$ with high probability. Yet no efficient algorithm is know to find a satisfying assignment for densities as low as $m/n \sim r_k \cdot \ln(k)/k$ with a non-vanishing probability. In fact, the density $m/n \sim r_k \cdot \ln(k)/k$ seems to form a barrier for a broad class of local search algorithms. One of the very few algorithms that plausibly seemed capable of breaking this barrier is a message passing algorithm called *belief propagation guided decimation*. It was put forward on the basis of deep but non-rigorous statistical mechanics considerations. Experiments conducted for $k = 3, 4, 5$ suggested that the algorithm might succeed for densities very close to r_k. Furnishing the first rigorous analysis of BP decimation, the present paper shows that the algorithm *fails* to find a satisfying assignment already for $m/n \geq \rho \cdot r_k/k$, for a constant $\rho > 0$ independent of k.

Key words: random discrete structures, k-SAT, message passing algorithms, belief propagation

---

*Supported by EPSRC grant EP/G039070/2 and DIMAP.

1 Introduction and results

Let $k \geq 3$ and $n > 1$ be integers, let $r > 0$ be a real, and set $m = \lceil rn \rceil$. Let $\mathbf{\Phi} = \mathbf{\Phi}_k(n, m)$ be a propositional formula obtained by choosing a set of m clauses of length k over the variables $x_1, \ldots, x_n$ uniformly at random. For k, r fixed we say that $\mathbf{\Phi}$ has some property $\mathcal{P}$ *with high probability* ('w.h.p.') if $\lim_{n\to\infty} \mathrm{P}\left[\mathbf{\Phi} \in \mathcal{P}\right] = 1$.

Non-constructive arguments show that w.h.p. $\mathbf{\Phi}$ is satisfiable if $r < r_k = 2^k \ln 2 - k$, and unsatisfiable if $r > 2^k \ln 2$ [3, 5]. However, in spite of several decades of extensive research in the CS community, both theoretical and experimental, no algorithm has emerged to find a satisfying assignment for densities r anywhere close to r_k in polynomial time w.h.p. More precisely, the best current polynomial time algorithm, designed specifically to solve random k-CNFs, fails for $r > r_k \cdot \frac{\ln k}{k}$ w.h.p. [12]. Furthermore, a great many algorithms are known to fail for asymptotically even smaller densities $r > r_k \cdot \frac{\rho}{k}$, where $\rho > 0$ is a constant (independent of k) that depends on the sophistication of the algorithm. Examples include simple linear-time algorithms such as Unit Clause ($\rho = \mathrm{e}/2$) [10] or Shortest Clause ($\rho = 1.817$) [15], as well as a wide range of DPLL-type algorithms ($\rho = 11/4$) [1]. More generally, in [2] we provided evidence that the density $r = r_k \cdot \frac{\ln k}{k}$ may form a barrier for a broad class of local search algorithms.

Against the background of these negative results, it came as a surprise when experiments indicated that certain *message passing algorithms* solve random k-CNFs in polynomial time for densities r very close to the conjectured thresholds for the existence of satisfying assignments, at least for $k = 3, 4, 5$ [8, 19, 21]. These algorithms, namely *Belief Propagation guided decimation* and *Survey Propagation guided decimation*, were put forward on the basis of very insightful but highly non-rigorous statistical mechanics considerations [8, 20]. Indeed, these message passing algorithms have come to be considered the only plausible 'candidates' for overcoming the $r_k \cdot \frac{\ln k}{k}$ barrier from [2]. Therefore, since these algorithms were suggested, one of the most important challenges in this area has been to come up with a rigorous analysis [4].

The present paper contributes the first rigorous analysis of Belief Propagation guided decimation for random k-SAT and, in fact, for any random constraint satisfaction problem. We establish a *negative* result: for general k the BP decimation algorithm `BPdec` fails to find a satisfying assignment w.h.p. for densities $r > r_k \cdot \frac{\rho}{k}$ for a certain constant $\rho > 0$. In other words, for general k `BPdec` does not outperform other, simpler algorithms (such as the one from [12]).

Stating the result precisely requires a little care, because it involves two levels of randomness: the choice of the random formula $\mathbf{\Phi}$, and the 'coin tosses' of the randomized algorithm `BPdec`. For a (fixed, non-random) k-CNF Φ let $\mathrm{success}(\Phi)$ denote the probability that `BPdec`(Φ) outputs a satisfying assignment (where, of course, 'probability' refers to the coin tosses of the algorithm). Then, if we apply `BPdec` to the *random* k-CNF $\mathbf{\Phi}$, the success probability $\mathrm{success}(\mathbf{\Phi})$ becomes a random variable. Recall that $\mathbf{\Phi}$ is unsatisfiable for $r > 2^k \ln 2$ w.h.p.

Theorem 1.1 *There is a constant $\rho_0 > 0$ such that for any k, r satisfying*

$$\rho_0 \cdot 2^k/k \leq r \leq 2^k \ln 2 \tag{1}$$

we have $\mathrm{success}(\mathbf{\Phi}) \leq \exp(-\Omega(n))$ *w.h.p.*

Theorem 1.1 contrasts with the rather promising experimental results on the performance of `BPdec`$(\mathbf{\Phi})$. But these results were obtained for 'small' $k = 3, 4, 5$. Already for $k = 10$ experiments are difficult to conduct, because the relevant density r scales exponentially in k. On the other hand, the constant ρ_0 in Theorem 1.1 implies a lower bound on k, as for small k the range (1) may just be empty. (Since the analysis is intricate as is, no attempt has been made to compute or optimize the precise value of ρ_0.)

An objection to a negative result like Theorem 1.1 could be that it might hinge on a small detail of the algorithm that could easily be fixed. However, in the sequel we will see that our analysis implies that in the regime (1) the key hypothesis on which `BPdec` is based is untrue.

The Belief Propagation guided decimation algorithm. [1] Fix a satisfiable k-CNF Φ on the variables $V = \{x_1, \ldots, x_n\}$. The algorithm `BPdec` is an attempt at implementing the following thought experiment, which yields an assignment $\sigma : V \to \{0, 1\}$ (with 0 representing 'false' and 1 representing 'true').

Experiment 1.2 Input: *A satisfiable k-CNF Φ.* Result: *An assignment $\sigma : V \to \{0, 1\}$.*

---

[1] In addition to the following discussion, more details on Belief/Survey propagation and the statistical mechanics perspective on random k-SAT can be found in Appendix A.

```
0. Let $\Phi_0 = \Phi$.
1. For $t = 1, \ldots, n$ do
2.     Compute the fraction $M_{x_t}(\Phi_{t-1})$ of all satisfying assignments of $\Phi_{t-1}$ in which the variable $x_t$ takes the value $1$.
3.     Assign $\sigma(x_t) = 1$ with probability $M_{x_t}(\Phi_{t-1})$, and let $\sigma(x_t) = 0$ otherwise.
4.     Obtain the formula $\Phi_t$ from $\Phi_{t-1}$ by substituting the value $\sigma(x_t)$ for $x_t$ and simplifying (i.e., delete all clauses that got satisfied, and omit $x_t$ from all other clauses).
5. Return the assignment $\sigma$.
```

A moment's reflection reveals that the above experiment not only produces a satisfying assignment, but that its (random) outcome is in fact uniformly distributed over the set of all satisfying assignments of Φ. The obvious obstacle to actually implementing this experiment is the computation of the marginal probability $M_{x_t}(\Phi_{t-1})$ that x_t takes the value 'true' in a random satisfying assignment of Φ_{t-1}. Indeed, computing $M_{x_t}(\Phi_{t-1})$ is $\#P$-hard in the worst case.

Yet, under what conditions could we hope to compute (or 'approximate') the marginals $M_x(\Phi_{t-1})$ in the 'decimated' formula Φ_{t-1}? Clearly, the $M_x(\Phi_{t-1})$ are influenced by 'local' effects. For instance, if x occurs in a *unit clause* a of Φ_{t-1}, i.e., a clause whose other $k-1$ variables have been assigned already without satisfying a, then x *must* be assigned so as to satisfy a. Hence, if x appears in a positively, then $M_x(\Phi_{t-1}) = 1$, and otherwise $M_x(\Phi_{t-1}) = 0$. Similarly, if x occurs *only* positively in Φ_{t-1}, then $M_x(\Phi_{t-1}) \geq 1/2$. More intricately, if x occurs in a clause a that contains another variable y that appears is a unit clause b, then this will affect the marginal of x.

The key hypothesis underlying `BPdec` is that in random formulas such local effects *determine* the marginals $M_x(\Phi_{t-1})$ asymptotically. To define 'local' precisely, we need a metric on the variables/clauses. This metric is the one induced by the *factor graph* $G = G(\Phi_{t-1})$ of Φ_{t-1}, which is a bipartite graph whose vertices are the variables $V_t = \{x_t, \ldots, x_n\}$ and the clauses of Φ_{t-1}. Each clause is adjacent to the variables that occur in it. For an integer $\omega \geq 1$ let $N^\omega(x_t)$ signify the set of all vertices of G that have distance at most 2ω from x_t. Then the induced subgraph $G\left[N^\omega(x_t)\right]$ corresponds to the sub-formula of Φ_{t-1} obtained by removing all clauses and variables at distance more than 2ω from x_t. Note that all vertices at distance precisely 2ω are variables, so that any satisfying assignment of Φ induces a satisfying assignment of the sub-formula. Let us denote by $M_{x_t}(\Phi_{t-1}, \omega)$ the marginal probability that x_t takes the value 1 in a random satisfying assignment of this sub-formula.

Of course, for a 'worst-case' formula Φ the 'local' marginals $M_{x_t}(\Phi_{t-1}, \omega)$ may be just as difficult to compute as the overall marginals $M_{x_t}(\Phi_{t-1})$ themselves. Therefore, `BPdec` employs an efficient heuristic called *Belief Propagation* ('BP'), which yields certain values $\mu_{x_t}(\Phi_{t-1}, \omega) \in [0,1]$; we will state this heuristic below. If the induced subgraph $G\left[N^\omega(x_t)\right]$ is a tree, then indeed $\mu_{x_t}(\Phi_{t-1}, \omega) = M_{x_t}(\Phi_{t-1}, \omega)$. Moreover, standard arguments show that in a random formula $\mathbf{\Phi}$ actually $G\left[N^\omega(x_t)\right]$ is a tree w.h.p. so long as $\omega = o(\ln n)$. More generally, in order to obtain an efficient algorithm it would be sufficient for the BP outcomes $\mu_{x_t}(\Phi_{t-1}, \omega)$ to approximate the true overall marginals $M_{x_t}(\Phi_{t-1})$ well for *some* (say, polynomially computable, polynomially bounded) function $\omega = \omega(n) \geq 1$. This leads to the following hypothesis underpinning `BPdec` (cf. [20]).

Hypothesis 1.3 *With probability $1 - o(1)$ over the choice of $\mathbf{\Phi}$ and the random decisions in Experiment 1.2 the following holds for all $1 \leq t \leq n$.*

- *i. For any $\varepsilon > 0$ there is $\omega = \omega(\varepsilon, k, r)$ such that $|M_{x_t}(\mathbf{\Phi}_{t-1}) - M_{x_t}(\mathbf{\Phi}_{t-1}, \omega)| \leq \varepsilon$.*
- *ii. There is a function $\omega = \omega(n) \geq 1$ such that $M_{x_t}(\mathbf{\Phi}_{t-1}) = \mu_{x_t}(\mathbf{\Phi}_{t-1}, \omega) + o(1)$ w.h.p.*

Hypothesis 1.3 motivates the following algorithm [24], which is called *BP guided decimation* because it combines BP (Step 2) with a procedure to 'decimate' the formula (Steps 3 and 4).[2]

Algorithm 1.4 `BPdec`(Φ)
Input: A k-CNF Φ on $V = \{x_1, \ldots, x_n\}$. *Output:* An assignment $\sigma : V \to \{0,1\}$.

```
0. Let $\Phi_0 = \Phi$.
1. For $t = 1, \ldots, n$ do
2.     Use BP to compute $\mu_{x_t}(\Phi_{t-1}, \omega)$.
3.     Assign $\sigma(x_t) = 1$ with probability $\mu_{x_t}(\Phi_{t-1}, \omega)$, and let $\sigma(x_t) = 0$ otherwise.
```

---

[2] `BPdec` is the simplest, but arguably also the most natural version of a 'BP decimation' algorithm. Appendix A contains a discussion of other version of BP decimation.

4. Obtain the formula Φ_t from Φ_{t-1} by substituting the value $\sigma(x_t)$ for x_t and simplifying.
5. Return the assignment σ.

Remark 1.5 *The statement of Theorem 1.1 holds* regardless *of the function $\omega = \omega(n) \geq 0$ chosen in Step 2.*

Although, strictly speaking, Hypothesis 1.3 provides neither a necessary nor a sufficient condition for `BPdec` to succeed on random k-CNFs w.h.p., the hypothesis inspired the algorithm. In fact, it plays an important role in the statistical mechanics work on random k-SAT generally, see Section 2 below. The proof of Theorem 1.1 entails

Corollary 1.6 *Both statements of Hypothesis 1.3 are false for k, r satisfying (1).*

We finish this section by briefly defining Belief Propagation for k-SAT; for a detailed derivation we refer the reader to [8, 26]. Ultimately, we need to define the value $\mu_{x_t}(\Phi_{t-1}, \omega)$ for Step 2 of `BPdec`.

Let $N(v)$ denote the neighborhood of a vertex v of the factor graph $G(\Phi_{t-1})$. For a variable $x \in V_t$ and a clause $a \in N(x)$ we will denote the ordered pair (x, a) by $x \to a$. Similarly, $a \to x$ stands for the pair (a, x). Furthermore, we let $\mathrm{sign}(x, a) = 1$ if x occurs in a positively, and $\mathrm{sign}(x, a) = -1$ otherwise.

The *message space* $M(\Phi_{t-1})$ is the set of all tuples $(\mu_{x\to a}(\zeta))_{x\in V_t,\, a\in N(x),\, \zeta\in\{0,1\}}$ such that $\mu_{x\to a}(\zeta) \in [0, 1]$ and $\mu_{x\to a}(0) + \mu_{x\to a}(1) = 1$ for all x, a, ζ. For $\mu \in M(\Phi)$ we define

$$\mu_{a\to x}(\zeta) \quad = \quad \begin{cases} 1 & \text{if } \zeta = (1 + \mathrm{sign}(x, a))/2, \\ 1 - \prod_{y\in N(a)\setminus\{x\}} \mu_{y\to a}\left(\frac{1-\mathrm{sign}(y,a)}{2}\right) & \text{otherwise.} \end{cases} \tag{2}$$

Furthermore, we define the ***belief propagation operator*** BP as follows: for any $\mu \in M(\Phi_{t-1})$ we define $\mathrm{BP}(\mu) \in M(\Phi_{t-1})$ by letting

$$(\mathrm{BP}(\mu))_{x\to a}(\zeta) \quad = \quad \frac{\prod_{b\in N(x)\setminus\{a\}} \mu_{b\to x}(\zeta)}{\prod_{b\in N(x)\setminus\{a\}} \mu_{b\to x}(0) + \prod_{b\in N(x)\setminus\{a\}} \mu_{b\to x}(1)} \tag{3}$$

unless the denominator equals zero, in which case $(\mathrm{BP}(\mu))_{x\to a}(\zeta) = \frac{1}{2}$.

Finally, the values $\mu_x(\Phi_t, \omega)$ that are used in step 3 of `BPdec` are defined as follows. Let $\mu\,[0] = \frac{1}{2} \cdot \mathbf{1} \in M(\Phi_{t-1})$ be the vector with all entries equal to $\frac{1}{2}$. Moreover, define inductively $\mu\,[\ell] = \mathrm{BP}(\mu\,[\ell - 1])$ for $1 \leq \ell \leq \omega$. Then

$$\mu_x(\Phi_{t-1}, \omega) = \frac{\prod_{b\in N(x)} \mu_{b\to x}(1)\,[\omega]}{\prod_{b\in N(x)} \mu_{b\to x}(0)\,[\omega] + \prod_{b\in N(x)} \mu_{b\to x}(1)\,[\omega]} \quad \text{for any } x \in V_t, \tag{4}$$

unless the denominator is zero, in which case we set $\mu_x(\Phi_{t-1}, \omega) = \frac{1}{2}$.

The intuition here is that the $\mu_{x\to a}(\zeta)$ are ‘messages’ from a variable x to the clauses a in which x occurs, indicating how likely x were to take the value ζ if clause a were removed from the formula. Based on these, (2) yields messages $\mu_{a\to x}(\zeta)$ from clauses a to variables x, indicating the probability that a is satisfied if x takes the value ζ and all other variables $y \in N(a) \setminus \{x\}$ are assigned independently to either value $\xi \in \{0, 1\}$ with probability $\mu_{y\to a}(\xi)$. The BP operator (3) then uses these messages $\mu_{a\to x}$ in order to ‘update’ the messages from variables to clauses. More precisely, for each x and $a \in N(x)$ the new messages $(\mathrm{BP}(\mu))_{x\to a}(\zeta)$ are computed under the hypothesis that all other clauses $b \in N(x) \setminus \{a\}$ are satisfied with probabilities $\mu_{b\to x}(\zeta)$ independently if x takes the value ζ. Finally, the difference between (3) and (4) is that the latter product runs over *all* clauses $b \in N(x)$. An inductive proof shows that, if for a variable x the subgraph $G\,[N^\omega(x)]$ of the factor graph is a tree, then in fact $\mu_x(\Phi_t, \omega) = M_x(\Phi_t, \omega)$ [8].

2 Related work

2.1 The statistical mechanics perspective.

The statistical mechanics work on random k-SAT is based on the ‘cavity method’, an ingenious, albeit highly non-rigorous, technique from the theory of spin glasses [21]. This actually is an *analytic* technique that yields, among other things, a conjecture on the precise location of the thresholds for satisfiability for any k. `BPdec` can be viewed as an attempt at turning (a part of) the cavity method into an efficient algorithm [8, 20, 24].

However, even the cavity method does not yield direct, quantitative estimates as to the regime of clause densities r for which `BPdec`$(\mathbf{\Phi})$ succeeds w.h.p. The obstacle is the decimation step of `BPdec`, i.e., the fact that `BPdec` eliminates one variable in each iteration, and then repeats the Belief Propagation computation on the reduced formula. In fact, tracing the BP decimation process directly (which is the approach taken in the present paper) has been attributed as "extremely challenging" [24].

In an important paper [20] the cavity method was used to determine the density r_c up to which statements i. and ii. in Hypothesis 1.3 hold for $t = 1$, i.e., in the *undecimated* random formula $\mathbf{\Phi}$. This density r_c is called the *condensation point*. The cavity method yields $r_c = 2^k \ln 2 - \frac{3}{2}\ln 2$, which is very close to the conjectured satisfiability threshold $2^k \ln 2 - \frac{1}{2}(1 + \ln 2)$. Indeed, [20] goes on to forecast that `BPdec` "should perform well up to the condensation point" r_c.

The issue with this prognosis (which would not sit well with Theorem 1.1) is that it disregards the effect of the decimation process, as pointed out recently in [27]. That paper extends the cavity method to analyze Experiment 1.2, i.e., the idealized procedure that `BPdec` was designed to implement. This (non-rigorous) analysis suggests that by decimating, i.e., by assigning and eliminating variables, for t sufficiently close to n the reduced formula $\mathbf{\Phi}_t$ exhibits a 'condensation' phenomenon even at densities r far below r_c. Quantitatively, the analysis from [27] suggests that Hypothesis 1.3 is invalid for $r > \mathrm{e} \cdot 2^k/k$. Yet Hypothesis 1.3 is neither necessary nor sufficient for the success of `BPdec`$(\mathbf{\Phi})$. Indeed, the paper [27], which supersedes an earlier non-rigorous study [24], mentions a rigorous analysis of `BPdec` explicitly as an open problem.

2.2 Rigorous work.

Belief Propagation (i.e., the message passing procedure to 'approximate' the marginals of a probability distribution described by a factor graph) is a fairly generic technique [26]. If the factor graph is acyclic, then Belief Propagation provably computes the correct marginals, which is generally not true in the presence of cycles. Nevertheless, BP is often applied as a heuristic to factor graphs with cycles anyway in the hope that it may yield useful approximations.

In spite of BP's practical success (and popularity), rigorous analyses are scarce. A few exist in the context of LDPC decoding [22, 28]. We also analyzed BP for graph 3-coloring [13] on a certain class of expander graphs. A further related result deals with another, conceptually far simpler message passing algorithm called Warning Propagation ('WP') on certain random 3-CNFs ("planted model") [14].

The big conceptual difference between the present and the previous work is that the algorithm `BPdec` studied here combines BP with *decimation* (i.e., the algorithm assigns one variable at a time and reruns BP on the reduced formula). The present paper furnishes the first analysis of this kind of algorithm. By comparison, previous analyses deal with algorithms that use BP in a 'one shot' fashion, i.e., the 'marginals' obtained via BP are used directly to assign all variables *at once* [13, 22, 28]. Roughly speaking, this approach works when the problem instances are somewhat over-constrained so that there is (essentially) a unique solution. By contrast, random k-CNFs with densities $r < r_k$, i.e., below the k-SAT threshold have *exponentially* many satisfying assignments whose typical pairwise distance is close to $n/2$ w.h.p. [5].

The study of the BP marginals on the undecimated random formula $\mathbf{\Phi}$ is somewhat related to the so-called *reconstruction problem*. This problem has been studied on 'symmetric' random CSPs, which include problems such as (hyper)graph coloring, but not k-SAT [23]. The proofs are based on indirect arguments (closely related to the second moment method), which do not seem to extend to an analysis of `BPdec`.

Finally, the paper [25] establishes $r \sim \frac{2\ln k}{k}$ as the threshold for *Gibbs uniqueness* in random k-SAT. The proof is via studying the correlation decay of formulas whose factor graphs are random trees. The result implies that up to density $r \sim \frac{2\ln k}{k}$ it is true that $M_x(\mathbf{\Phi}) = \mu_x(\mathbf{\Phi}, \omega) + o(1)$ if $1 \ll \omega \ll \ln n$ w.h.p. That is, BP does actually approximate the true marginals well on the original, undecimated formula. However, this result does not extend to an analysis of the decimation process either. Note that the density $r \sim \frac{2\ln k}{k}$ to which [25] applies is far below the $\Omega(2^k/k)$ addressed in Theorem 1.1.

3 The probabilistic framework for analyzing `BPdec`

The single most important technique for analyzing algorithms on the random k-CNF $\mathbf{\Phi}$ is the *method of deferred decisions*. Where it applies, the dynamics of the algorithm can typically be traced tightly via differential equations, martingales, or Markov chains. Virtually all of the previous analyses of algorithms for random k-SAT are based on this approach [6, 9, 10, 11, 12, 15, 16, 18]. Unfortunately, the 'deferred decisions' technique is limited to very simple, 'shortsighted' algorithms that decide upon the value of a variable x on the basis of the clauses/variables at distance, say, one or two from x in the factor graph [6]. By contrast, in order to assign the tth variable x_t, `BPdec` explores clauses at distance up to 2ω from x_t, where (potentially) $\omega = \omega(n) \to \infty$. This renders a 'deferred decisions' approach hopeless.

Therefore, to prove Theorem 1.1 we need a fundamentally different strategy. In the present section we set up the probabilistic framework for the analysis. We will basically reduce the analysis of `BPdec` to the problem of analyzing the BP operator on the formula that is obtained from the random $\mathbf{\Phi}$ by substituting 'true' for the first t variables $x_1, \ldots, x_t$ (Theorem 3.3 blow). In the next section we will show that this decimated formula enjoys a few simple quasirandomness properties with probability extremely close to one. Finally, in Section 5 we will show that these quasirandomness properties suffice in order to trace the dynamics of the BP computation.

Applied to a k-CNF Φ on $V = \{x_1, \ldots, x_n\}$, `BPdec` yields an assignment $\sigma : V \to \{0,1\}$ (that may or may not be satisfying). This assignment is random (even if Φ is not), because `BPdec` itself is randomized. Hence, for any fixed Φ running `BPdec`(Φ) induces a probability distribution β_Φ on $\{0,1\}^V$. Let $S(\Phi)$ denote the set of all satisfying assignments of Φ. Then the 'success probability' of `BPdec` on Φ is just $\mathrm{success}(\Phi) = \beta_\Phi(S(\Phi))$. Thus, to establish Theorem 1.1 we need to show that in random formulas $\mathrm{success}(\mathbf{\Phi}) = \beta_{\mathbf{\Phi}}(S(\mathbf{\Phi}))$ is exponentially small w.h.p. To this end, we are going to prove that the measure $\beta_{\mathbf{\Phi}}$ is 'close' to the uniform distribution on $\{0,1\}^V$ w.h.p., of which $S(\mathbf{\Phi})$ constitutes only an exponentially small fraction.

For the sake of the analysis, we consider a slightly modified version of `BPdec`. While the original `BPdec` assigns the variables in the natural order $x_1, \ldots, x_n$, the modified version `PermBPdec` chooses a permutation π of $[n]$ uniformly at random and assigns the variables in the order $x_{\pi(1)}, \ldots, x_{\pi(n)}$. Let $\bar\beta_\Phi$ denote the probability distribution induced on $\{0,1\}^V$ by `PermBPdec`(Φ). Because the uniform distribution over k-CNFs is invariant under permutations of the variables, we obtain

Fact 3.1 *If* $\bar\beta_{\mathbf{\Phi}}(S(\mathbf{\Phi})) \leq \exp(-\Omega(n))$ *w.h.p., then* $\mathrm{success}(\mathbf{\Phi}) = \beta_{\mathbf{\Phi}}(S(\mathbf{\Phi})) \leq \exp(-\Omega(n))$ *w.h.p.*

Let Φ be a k-CNF and let $\delta > 0$. Given a permutation π and a partial assignment $\sigma : \left\{x_{\pi(s)} : s \leq t\right\} \to \{0,1\}$, we let $\Phi_{t,\pi,\sigma}$ denote the formula obtained from Φ by substituting the values $\sigma(x_{\pi(s)})$ for the variables $x_{\pi(s)}$ for $1 \leq s \leq t$ and simplifying. For a number $\delta > 0$ and an index $l > t$ we say that $x_{\pi(l)}$ is (δ,t)-***biased*** if the result $\mu_{x_{\pi(l)}}(\Phi_{t,\pi,\sigma}, \omega)$ of the BP computation on $\Phi_{t,\pi,\sigma}$ differs from $\frac{1}{2}$ by more than δ, i.e.,

$$\left|\mu_{x_{\pi(l)}}(\Phi_{t,\pi,\sigma}, \omega) - 1/2\right| > \delta.$$

Moreover, the triple (Φ, π, σ) is (δ,t)-***balanced*** if no more than $\delta(n-t)$ variables are (δ,t)-biased.

Let π be the permutation chosen by `PermBPdec`(Φ), and let σ be the partial assignment constructed in the first t steps. The variable $x_{\pi(t+1)}$ is uniformly distributed over the set $V \setminus \left\{x_{\pi(s)} : s \leq t\right\}$ of currently unassigned variables. Hence, if (Φ, π, σ) is (δ,t)-balanced, then the probability that $x_{\pi(t+1)}$ is (δ,t)-biased is bounded by δ. (This conclusion was the purpose of decimating the variables in a random order.) Furthermore, given that $x_{\pi(t+1)}$ is not (δ,t)-biased, the probability that `PermBPdec` will assigned it the value 'true' lies in the interval $\left[\frac{1}{2} - \delta, \frac{1}{2} + \delta\right]$. Consequently,

$$\left|\frac{1}{2} - \mathrm{P}\left[\sigma(x_{\pi(t+1)}) = 1 | (\Phi, \pi, \sigma) \text{ is } (\delta,t)\text{-balanced}\right]\right| \leq 2\delta. \tag{5}$$

Thus, the smaller δ, the closer $\sigma(x_{\pi(t+1)})$ comes to being uniformly distributed. Hence, if (δ,t)-balancedness holds for all t with a 'small' δ, then $\bar\beta_\Phi$ will be close to the uniform distribution on $\{0,1\}^V$.

To put this observation to work, we define

$$\delta_t = \exp(-c(1-t/n)k) \text{ and } T = (1 - r/2^k)n, \tag{6}$$

where $c > 0$ is a small absolute constant (independent of k, r, t, n). Furthermore, for $\xi > 0$ we say that Φ is (t,ξ)-***uniform*** if the total number of permutations π and assignments $\sigma \in \{0,1\}^V$ such that (Φ, π, σ) is not (δ_t, t)-balanced is bounded by $2^n n! \cdot \exp\left[-\xi n - 10\sum_{s=1}^t \delta_s\right]$. Proceeding by induction on t, we can use (5) to relate the distribution $\bar\beta_\Phi$ to the uniform distribution on $\{0,1\}^V$ for formulas that are (t,ξ)-uniform.

Proposition 3.2 *Suppose that* Φ *is* (t,ξ)*-uniform for all* $0 \leq t \leq T$. *Then*

$$\bar\beta_\Phi(\mathcal{E}) \leq \frac{|\mathcal{E}|}{2^T} \cdot \exp\left[4\sum_{t=1}^T \delta_t\right] + \exp(-\xi n/2) \quad \textit{for any } \mathcal{E} \subset \{0,1\}^V. \tag{7}$$

Proposition 3.2 reduces the proof of Theorem 1.1 to showing that $\mathbf{\Phi}$ is (t,ξ)-uniform. The following theorem provides the key step.

Theorem 3.3 *There is a constant* $\rho_0 > 0$ *such that for any* k, r *satisfying* $2^k\rho_0/k \leq r \leq 2^k \ln 2$ *there is* $\xi = \xi(k,r) > 0$ *so that for* n *large enough the following holds. Fix any permutation* π *of* $[n]$ *and any assignment* $\sigma \in \{0,1\}^V$. *Then for any* $0 \leq t \leq T$ *we have*

$$\mathrm{P}\left[(\mathbf{\Phi},\pi,\sigma) \textit{ is } (\delta_t,t)\textit{-balanced}\right] \geq 1 - \exp\left[-2\xi n - 10\sum_{s=1}^{t}\delta_s\right]. \tag{8}$$

Theorem 3.3 shows that for a *fixed* pair (π,σ) the triple $(\mathbf{\Phi},\pi,\sigma)$ is likely to be balanced. A simple first moment argument then yields a bound on the *number* of pairs (π,σ) so that $(\mathbf{\Phi},\pi,\sigma)$ fails to be balanced. This implies that $\mathbf{\Phi}$ is likely to be (t,ξ)-uniform.

Corollary 3.4 *In the notation of Theorem 3.3,* $\mathrm{P}\left[\forall t \leq T : \mathbf{\Phi} \textit{ is } (t,\xi)\textit{-uniform}\right] \geq 1 - \exp(-\xi n)$.

Combining Proposition 3.2 and Corollary 3.4, we see that (7) holds for a random formula $\mathbf{\Phi}$ w.h.p. In particular, we can apply (7) to the set $\mathcal{E} = S(\mathbf{\Phi})$ of satisfying assignments. A simple first moment calculation shows that $|\mathcal{E}| \leq 2^{(1+o(1))n}(1-2^{-k})^m \leq \exp[n(\ln 2 - r2^{-k} + o(1))]$ w.h.p. Plugging this bound into (7), recalling that $T = (1 - r/2^k)n$, and observing $\sum_{t=1}^{T}\delta_t \sim n/(ck\exp(ckr/2^k))$ yields Theorem 1.1.

4 The quasirandomness property

The goal in this section is to establish Theorem 3.3, which states that for any fixed permutation π and any fixed assignment σ the triple $(\mathbf{\Phi},\pi,\sigma)$ is (δ_t,t)-balanced with probability very close to one. By the symmetry properties of the uniformly distribution over k-CNFs, we may assume without loss of generality that $\pi = \mathrm{id}$ is the identity permutation, and that $\sigma = \mathbf{1}$ is the all-true assignment. Then the decimated formula $\mathbf{\Phi}_t = \mathbf{\Phi}_{t,\pi,\sigma}$ is obtained from $\mathbf{\Phi}$ just by substituting the value 'true' for $x_1,\ldots,x_t$ and simplifying. In this section we will exhibit a few simple quasirandomness properties that $\mathbf{\Phi}_t$ is very likely to possess. In Section 5 we will show that these properties suffice to trace the BP operator.

To state the quasirandomness properties, fix a k-CNF Φ. Let $\Phi_t = \Phi_{t,\mathrm{id},\mathbf{1}}$ denote the CNF obtained from Φ by substituting 'true' for $x_1,\ldots,x_t$ and simplifying ($1 \leq t \leq n$). Let $V_{t+1} = \{x_{t+1},\ldots,x_n\}$ be the set of variables of Φ_t. As before, we will denote the factor graph of Φ_t by $G = G(\Phi_t)$, and the neighborhood of a vertex v by $N(v)$. We call a clause a of Φ ***redundant*** if there is another clause b such that $|N(a)\cap N(b)| \geq 2$. Let

$$\theta = 1 - t/n$$

be the fraction of unassigned variables. Moreover, for a variable $x \in V_t$ and a set $Q \subset V_t$ let

$$N_{\leq 1}(x,Q) \quad = \quad \{b \in N(x) : |N(x)\cap Q\setminus\{x\}| \leq 1 \wedge 0.1\theta k \leq |N(b)| \leq 10\theta k\}. \tag{9}$$

Thus, $N_{\leq 1}(x,Q)$ is the set of all clauses that contain x (which may or may not be in Q) and at most one other variable from Q. In addition, there is a condition on the *length* $|N(b)|$ of the clause b in the decimated formula Φ_t. Observe that having assigned the first t variables, we should 'expect' the average clause length to be θk. For a linear map $\Lambda : \mathbf{R}^{V_t} \to \mathbf{R}^{V_t}$ let $\|\Lambda\|_\square$ signify the norm $\|\Lambda\|_\square = \max_{\zeta\in\mathbf{R}^{V_t}\setminus\{0\}} \frac{\|\Lambda\zeta\|_1}{\|\zeta\|_\infty}$.

Definition 4.1 *Let* $\delta > 0$. *We say that* Φ *is* ***(***δ,t***)-quasirandom*** *if* **Q0–Q4** *in Figure 1 are satisfied.*

Condition **Q0** simply bounds the number of redundant clauses and the number of variables of very high degree; it well-known to hold for random k-CNFs w.h.p. **Q1** provides a bound on the 'weight' of clauses in which variables $x \in V_{t+1}$ typically occur, where the weight of a clause b is $2^{-|N(b)|}$. Moreover, **Q2** imposes that there is no small set Q for which the total weight of the clauses touching that set is very big. In addition, **Q2** (essentially) requires that for most variables x the weights of the clauses where x occurs positively/negatively should approximately cancel out. **Q3** provides a bound on the lengths of clauses that contain many variables from a small set Q.

The most important condition is **Q4**. Think of a $V_{t+1}\times V_{t+1}$-matrix representation of the formula where the entry of a pair (x,y) of variables is the sum of the signed weights $2^{-N(b)}\mathrm{sign}(x,b)\mathrm{sign}(y,b)$ over clauses $b \in N(x)\cap N(y)$. This can be considered a weighted, signed adjacency matrix of Φ_t. Then **Q4** requires that this matrix has low discrepancy (i.e., it is 'flat'), modulo the clauses that contain more than one variable from some small 'exceptional' set Q.

Q0. *Φ has only $o(n)$ redundant clauses. Moreover, no more than $o(n)$ variables occur in more than $\ln n$ clauses of Φ.*

Q1. *No more than $10^{-5}\delta(n-t)$ variables occur in clauses of length less than $\theta k/10$ or greater than $10\theta k$. Moreover, there are at most $10^{-4}\delta(n-t)$ variables $x \in V_{t+1}$ such that*

$$(\theta k)^3 \delta \cdot \textstyle\sum_{b \in N(x)} 2^{-|N(b)|} > 1.$$

Q2. *If $Q \subset V_{t+1}$ has size $|Q| \leq \delta(n-t)$, then there are no more than $10^{-4}\delta(n-t)$ variables x such that either*

$$\sum_{b \in N(x) : |N(b) \cap Q \setminus \{x\}| = 1} 2^{-|N(b)|} > (\theta k)^5 \delta, \quad \textit{or} \qquad \sum_{b \in N(x) : |N(b) \cap Q \setminus \{x\}| > 1} 2^{|N(b) \cap Q \setminus \{x\}| - |N(b)|} > \frac{\delta}{\theta k}, \textit{ or}$$

$$\left| \sum_{b \in N_{\leq 1}(x, Q)} \frac{\mathrm{sign}(x,b)}{2^{|N(b)|}} \right| > \frac{\delta}{1000}.$$

Q3. *For any $0.01 \leq z \leq 1$ and any set $Q \subset V_{t+1}$ of size $0.01\delta(n-t) \leq |Q| \leq 100\delta(n-t)$ we have*

$$\sum_{b : |N(b) \cap Q| \geq z|N(b)|} |N(b)| \leq 1.01|Q|/z.$$

Q4. *For any set $Q \subset V_{t+1}$ of size $|Q| \leq 10\delta(n-t)$ the linear operator*

$$\Lambda_Q : \mathbf{R}^{V_{t+1}} \to \mathbf{R}^{V_{t+1}}, \qquad \Gamma \mapsto \left[\sum_{b \in N_{\leq 1}(x, Q)} \sum_{y \in N(b) \setminus \{x\}} 2^{-|N(b)|} \mathrm{sign}(x,b) \mathrm{sign}(y,b) \Gamma_y \right]_{x \in V_{t+1}}$$

has norm $\|\Lambda_Q\|_{\square} \leq \delta^4(n-t)$.

Figure 1: The conditions for Definition 4.1.

Proposition 4.2 *There exists a constant $\rho_0 > 0$ such that for any k, r satisfying $\rho_0 \cdot 2^k/k \leq r \leq 2^k \ln 2$ there is $\xi = \xi(k,r) > 0$ so that for n large and δ_t, T as in (6) we have*

$$\mathrm{P}\left[\mathbf{\Phi} \textit{ is } (\delta_t, t)\textit{-quasirandom}\right] \geq 1 - \exp\left[-10\left(\xi n + \textstyle\sum_{s \leq t} \delta_s\right)\right] \qquad \textit{for any } 1 \leq t \leq T.$$

Though non-trivial, the proof of Proposition 4.2 is based on standard arguments. Together with the following theorem, which we will establish in Section 5, Proposition 4.2 yields Theorem 3.3.

Theorem 4.3 *There is $\rho_0 > 0$ such that for any k, r satisfying $\rho_0 \cdot 2^k/k \leq r \leq 2^k \ln 2$ and n sufficiently large the following is true. Let Φ be a k-CNF with n variables and m clauses that is (δ_t, t)-quasirandom for some $1 \leq t \leq T = (1 - r/2^k)n$. Then $(\Phi, \mathrm{id}, \mathbf{1})$ is (δ_t, t)-balanced.*

5 Tracing the Belief Propagation operator

Throughout this section, we keep the notation from Section 4 and the assumptions of Theorem 4.3. Assuming that Φ is (δ_t, t)-quasirandom, we aim at showing that in the formula Φ_t obtained by substituting the value 'true' for the first t variables $x_1, \ldots, x_t$ the values $\mu_{x_s}(\Phi_t, \omega)$ computed by belief propagation are in the interval $\left[\frac{1}{2} - \delta_t, \frac{1}{2} + \delta_t\right]$, for all but at most $\delta_t(n-t)$ indices $t < s \leq n$.

To establish this, we are going to trace the BP operator on Φ_t, starting from the initial point $\mu[0] = \frac{1}{2} \cdot \mathbf{1}$. Let $\mu[\ell] = \mathrm{BP}^{\ell}(\mu[0]) \in M(\Phi)$ be the result of the first ℓ iterations of BP for $\ell \geq 1$. We say that a variable $x \in V_{t+1}$ is ***ℓ-biased*** if there is a clause $a \in N(x)$ such that $|\mu_{x \to a}(1)[\ell] - \frac{1}{2}| \geq 0.1\delta_t$. Let $B[\ell]$ signify the set of all ℓ-biased variables. The core of the proof is to establish the following proposition.

Proposition 5.1 *For all $\ell \geq 1$ we have $|B[\ell]| \leq \delta_t(n-t)$.*

Since the values $\mu_{x_s}(\Phi_t, \omega)$ are defined in terms of the 'messages' $\mu_{x \to a}(\zeta)[\omega]$ (cf. (4)), Proposition 5.1 (with $\ell = \omega$) and a straight computation yield that $(\Phi, \mathrm{id}, \mathbf{1})$ is (δ_t, t)-balanced, as claimed by Theorem 4.3.

The proof of Proposition 5.1 is by induction on ℓ. Roughly speaking, we need to verify that for 'most' variables x all messages $\mu_{x\to a}(1)\,[\ell+1]$ are close to $\frac{1}{2}$, assuming that $|B\,[\ell]| \le \delta_t(n-t)$. This is the same as saying that for most x the ratios

$$\frac{\mu_{x\to a}(1)\,[\ell+1]}{\mu_{x\to a}(0)\,[\ell+1]} = \prod_{b\in N(x)\setminus\{a\}} \frac{\mu_{b\to x}(1)\,[\ell]}{\mu_{b\to x}(0)\,[\ell]}. \tag{10}$$

are close to one, because by construction $\mu_{x\to a}(1)\,[\ell+1]+\mu_{x\to a}(0)\,[\ell+1]=1$, cf. (3). (One also has to consider the case that the denominator is zero, but this is a mere technicality.) To analyze the product (10), we decompose it into two pieces, the first of which, denoted by $P_{\le 1}(x\to a)$, comprises the factors for $b\in N_{\le 1}(x, B\,[\ell])\setminus\{a\}$ (cf. (9)). The following lemma summarizes the analysis of $P_{\le 1}(x\to a)$. (The analysis of the product $P_{>1}(x\to a)$ encompassing the remaining factors is omitted.)

Lemma 5.2 *For all but* $0.1\delta_t(n-t)$ *variables* $x\in V_{t+1}$ *we have* $\max_{a\in N(x)} |\ln P_{\le 1}(x\to a)| \le 0.01\delta_t$.

To prove Lemma 5.2, it is convenient to define $\Delta_{x\to a}=\mu_{x\to a}(1)\,[\ell]-\frac{1}{2}\in[-1/2,1/2]$. Moreover, we say that $x\in V_{t+1}$ is ***ℓ-good*** if it has the following three properties (cf. **Q0–Q2** in Figure 1).

G0. x does not occur in a redundant clause.

G1. $(\theta k)^3\delta_t\sum_{b\in N(x)}2^{-|N(b)|}\le 1$, and $0.1\theta k\le |N(b)|\le 10\theta k$ for all $b\in N(x)$.

G2. $|\sum_{b\in N_{\le 1}(b,B[\ell])}\mathrm{sign}(x,b)/2^{|N(b)|}|\le 0.001\delta_t$.

For ℓ-good variables x we are going to approximate $\ln P_{\le 1}(x\to a)$ by the *linear* expression

$$L_{x\to a} \quad = \sum_{b\in N_{\le 1}(x,B[\ell])\setminus\{a\}} 2^{1-|N(b)|}\left[\mathrm{sign}\,(x,b)+\sum_{y\in N(b)\setminus\{x\}}\mathrm{sign}\,(x,b)\,\mathrm{sign}\,(y,b)\,\Delta_{y\to b}\right], \tag{11}$$

which is nothing but the derivative of $-\ln P_{\le 1}(x\to a)$. A straight (if lengthy) computation yields

Lemma 5.3 *For any ℓ-good* $x\in V_{t+1}$ *and any* $a\in N(x)$ *we have* $|L_{x\to a}+\ln P_{\le 1}(x\to a)|\le 10^{-6}\delta_t$.

Proof of Lemma 5.2 (sketch). We can confine ourselves to ℓ-good variables, because conditions **Q0**, **Q1**, and **Q2** ensure the all but $0.05\delta_t(n-t)$ variables are ℓ-good. Due to Lemma 5.3, for ℓ-good variables it suffices to study the linear expression (11). For any ℓ-good x the term $\sum_{b\in N_{\le 1}(x,B[\ell])\setminus\{a\}}2^{1-|N(b)|}\mathrm{sign}\,(x,b)$ is bounded by $0.002\delta_t$ in absolute value by **G2**. Since this is the first bit of $L_{x\to a}$, we are left to bound

$$\xi_{x\to a}=\sum_{b\in N_{\le 1}(x,B[\ell])\setminus\{a\}}\ \sum_{y\in N(b)\setminus\{x\}}2^{1-|N(b)|}\mathrm{sign}\,(x,b)\,\mathrm{sign}\,(y,b)\,\Delta_{y\to b}. \tag{12}$$

for ℓ-good $x\in V_{t+1}$ and $a\in N(x)$.

The analysis of the $\xi_{x\to a}$ is via the linear operator $\Lambda_{B[\ell]}$ from condition **Q4**, whose definition clearly bears some resemblance to the r.h.s. of (12). The issue is that the dimension of the argument vector Δ of (12) is $\sum_{x\in V_{t+1}}|N(x)|$ (it has one component for each edge of the factor graph), whereas $\Lambda_{B[\ell]}$ only 'lives' in dimension $|V_{t+1}|=n-t$. Thus, to invoke **Q4** we will first have to project the vector Δ down to the lower dimensional space $\mathbf{R}^{V_{t+1}}$, and then show that this projection is not too distorting.

For any $y\in V_{t+1}$ let $\mathcal{N}(y)=\bigcup_{x\in V_{t+1}\setminus\{y\}}N(y)\cap N_{\le 1}(x,B\,[\ell])$ be the set of all clauses $b\in N(y)$ such that the entry $\Delta_{y\to b}$ contributes to some term $\xi_{x\to a}$. Then, for any $y\in V_{t+1}$ choose $b_y\in\mathcal{N}(y)$ arbitrarily and set $D_y=\Delta_{y\to b_y}\,[\ell]$ (and if $\mathcal{N}(y)=\emptyset$, then set $D_y=0$). The resulting 'projected' vector $D=(D_y)_{y\in V_{t+1}}$ has norm $\|D\|_\infty\le\frac{1}{2}$, because $|\Delta_{y\to b_y}\,[\ell]\,|\le\frac{1}{2}$ by definition. An elementary but lengthy argument shows that

$$|D_y-\Delta_{y\to b}|\le\delta_t^3 \quad \text{for any } b\in\mathcal{N}(y). \tag{13}$$

We can easily use **Q4** to analyze $\Lambda_{B[\ell]}D$: let

$$\zeta=(\zeta_x)_{x\in V_{t+1}}=\Lambda_{B[\ell]}D,\ \text{so that } \zeta_x=\sum_{b\in N_{\le 1}(x,B[\ell])}\ \sum_{y\in N(b)\setminus\{x\}}2^{1-|N(b)|}\mathrm{sign}(x,b)\mathrm{sign}(y,b)D_y.$$

As **Q4** ensures that $\left\|\Lambda_{B[\ell]}\right\|_{\Box} \leq \delta_t^4(n-t)$, we obtain $\|\zeta\|_1 = \left\|\Lambda_{B[\ell]} D\right\|_1 \leq \left\|\Lambda_{B[\ell]}\right\|_{\Box} \|D\|_\infty \leq \delta_t^4(n-t)$. Since $\|\zeta\|_1 = \sum_{x \in V_{t+1}} |\zeta_x|$, we see that

$$\left|\left\{x \in V_{t+1} : |\zeta_x| > \delta_t^2\right\}\right| \leq \delta_t^2(n-t). \tag{14}$$

To complete the proof, we need to relate ζ_x and $\xi_{x\to a}$ for $a \in N(x)$ for ℓ-good x. Just from the definitions of ζ_x and $\xi_{x\to a}$ we see that

$$\begin{aligned} |\xi_{x\to a} - \zeta_x| &\leq \sum_{y \in N(a)} 2^{1-|N(a)|} |D_y| + \sum_{b \in N_{\leq 1}(x,B[\ell]) \setminus \{a\}} \sum_{y \in N(b) \setminus \{x\}} 2^{1-|N(b)|} |\Delta_{y\to a} - D_y| \\ &\overset{(13)}{\leq} |N(a)| 2^{-|N(a)|} + \delta_t^3 \sum_{b \in N_{\leq 1}(x,B[\ell]) \setminus \{a\}} |N(b)| 2^{1-|N(b)|}. \end{aligned} \tag{15}$$

Since $|N(a)| \geq 0.1\theta k$ for all $a \in N(x)$ by **G1**, we get $|N(a)|2^{-|N(a)|} \leq \exp(-2c\theta k) = \delta_t^2 \leq \delta_t/10^4$, provided the constant $c > 0$ in the definition (6) of δ_t is chosen sufficiently small and $\theta k \geq kr/2^k \geq \rho_0$ for a sufficiently large constant $\rho_0 > 0$. Similarly, **G1** yields $\delta_t^3 \sum_{b \in N_{\leq 1}(x,B[\ell]) \setminus \{a\}} |N(b)| 2^{1-|N(b)|} \leq \delta_t/10^4$. Hence, $|\xi_{x\to a} - \zeta_x| \leq \delta_t/5000$ for any ℓ-good x. Thus, the assertion follows from (14). □

References

[1] D. Achlioptas, P. Beame, M. Molloy: Exponential bounds for DPLL below the satisfiability threshold. Proc. 15th SODA (2004) 139–140.

[2] D. Achlioptas, A. Coja-Oghlan: Algorithmic barriers from phase transitions. Proc. 49th FOCS (2008) 793–802.

[3] D. Achlioptas, C. Moore: Random k-SAT: two moments suffice to cross a sharp threshold. SIAM Journal on Computing **36** (2006) 740–762.

[4] D. Achlioptas, A. Naor, Y. Peres: Rigorous location of phase transitions in hard optimization problems. Nature **435** (2005) 759–764.

[5] D. Achlioptas, Y. Peres: The threshold for random k-SAT is $2^k \ln 2 - O(k)$. Journal of the AMS **17** (2004) 947–973.

[6] D. Achlioptas, G. Sorkin: Optimal myopic algorithms for random 3-SAT. Proc. 41st FOCS (2000) 590–600.

[7] J. Ardelius, L. Zdeborova: Exhaustive enumeration unveils clustering and freezing in random 3-SAT. Phys. Rev. E **78** (2008) 040101(R).

[8] A. Braunstein, M. Mézard, R. Zecchina: Survey propagation: an algorithm for satisfiability. Random Structures and Algorithms **27** (2005) 201–226.

[9] A. Broder, A. Frieze, E. Upfal: On the satisfiability and maximum satisfiability of random 3-CNF formulas. Proc. 4th SODA (1993) 322–330.

[10] M.-T. Chao, J. Franco: Probabilistic analysis of a generalization of the unit-clause literal selection heuristic for the k-satisfiability problem. Inform. Sci. **51** (1990) 289–314.

[11] V. Chvátal, B. Reed: Mick gets some (the odds are on his side). Proc. 33th FOCS (1992) 620–627.

[12] A. Coja-Oghlan: A better algorithm for random k-SAT. SIAM J. Computing **39** (2010) 2823–2864.

[13] A. Coja-Oghlan, E. Mossel, D. Vilenchik: A spectral approach to analyzing Belief Propagation for 3-coloring. Combinatorics, Probability and Computing **18** (2009) 881–912.

[14] U. Feige, E. Mossel, D. Vilenchik: Complete convergence of message passing algorithms for some satisfiability problems. Proc. 10th RANDOM (2006) 339–350.

[15] A. Frieze, S. Suen: Analysis of two simple heuristics on a random instance of k-SAT. Journal of Algorithms **20** (1996) 312–355.

[16] M. Hajiaghayi, G. Sorkin: The satisfiability threshold of random 3-SAT is at least 3.52. IBM Research Report RC22942 (2003).

[17] S. Janson, T. Łuczak, A. Ruciński: Random Graphs, Wiley 2000.

[18] A. Kaporis, L. Kirousis, E. Lalas: The probabilistic analysis of a greedy satisfiability algorithm. Random Structures and Algorithms **28** (2006) 444–480.

[19] L. Kroc, A. Sabharwal, B. Selman: Message-passing and local heuristics as decimation strategies for satisfiability. Proc 24th SAC (2009) 1408–1414.

[20] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, L. Zdeborova: Gibbs states and the set of solutions of random constraint satisfaction problems. Proc. National Academy of Sciences **104** (2007) 10318–10323.

[21] M. Mézard, G. Parisi, R. Zecchina: Analytic and algorithmic solution of random satisfiability problems. Science **297** (2002) 812–815.

[22] A. Montanari, R. Urbanke: Modern coding theory: the statistical mechanics and computer science point of view. Lecture notes, 2007.

[23] A. Montanari, R. Restrepo, P. Tetali: Reconstruction and clustering in random constraint satisfaction problems. Preprint (2009).

[24] A. Montanari, F. Ricci-Tersenghi, G. Semerjian: Solving constraint satisfaction problems through Belief Propagation-guided decimation. Proc. 45th Allerton (2007).

[25] A. Montanari, D. Shah: Counting good truth assignments of random k-SAT formulae. Proc. 18th SODA (2007) 1255–1264.

[26] J. Pearl: Probabilistic reasoning in intelligent systems: networks of plausible inference. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1988.

[27] F. Ricci-Tersenghi, G. Semerjian: On the cavity method for decimated random constraint satisfaction problems and the analysis of belief propagation guided decimation algorithms. J. Stat. Mech. (2009) P09001.

[28] T. Richardson, A. Shokrollahi, R. Urbanke: Design of capacity-approaching irregular low-density parity check codes. IEEE Trans. Info. Theory **47** (2001) 619–637.

Appendix

Appendix A contains a somewhat broader discussion of Belief and Survey Propagation and their statistical mechanics context, which has been omitted from the main body of the paper for lack of space.

The other appendices contain the complete proofs, as required by the submission guidelines. Appendix B deals with the proof of Proposition 3.2, and Appendix C contains the proof of Corollary 3.4. Then, in Appendix D we show how Theorem 1.1 follows from Proposition 3.2 and Corollary 3.4. The proof of Proposition 4.2 is given in Appendix E. Further, Appendix F gives a self-contained proof of Theorem 4.3 (it can be read independently of the outline given in Section 5). Finally, the proof of Corollary 1.6 is given in Appendix G.

A Message passing algorithms and statistical mechanics

A.1 Replica symmetry breaking and all that

Following the non-rigorous paper [20], we discuss in this section the statistical mechanics motivation for Hypothesis 1.3, which motivates `BPdec`. Let Φ be a k-CNF with variables $V = \{x_1, \ldots, x_n\}$. From a statistical mechanics perspective, the natural object of study is the set $S(\Phi)$ of all satisfying assignments of Φ, or, more precisely, the uniform probability distribution on this set. The uniform distribution on $S(\Phi)$ is called the *Gibbs measure* g_Φ.[3] For a *random* k-CNF $\mathbf{\Phi}$ the Gibbs measure $g_{\mathbf{\Phi}}$ becomes a random object, and the aim is to analyze the properties that it exhibits with probability $1 - o(1)$ over the choice of $\mathbf{\Phi}$.

As the clause/variable density r increases, the Gibbs measure $g_{\mathbf{\Phi}}$ undergoes several phase transitions. The first of these is called the *Gibbs uniqueness* transition. This is the point where the set $S(\mathbf{\Phi})$ of all satisftying assignments forms a connected subgraph of the Hamming cube $\{0,1\}^V$ w.h.p. The paper [25] establishes rigorously that this transition occurs at $r \sim r_u = \frac{2\ln k}{k}$. This result implies the two statements of Hypothesis 1.3 hold true for $r < r_u$ with $t = 1$, i.e., for the *undecimated* formula $\mathbf{\Phi}$. In fact, this is the maximum value of r for which this is (rigorously) known to be the case.

For densities $r > \frac{2\ln k}{k}$ but below the k-SAT threshold $r_s \sim 2^k \ln 2$ the subgraph of the Hamming cube induced on $S(\mathbf{\Phi})$ is expected to be disconnected. But for $r < r_d = 2^k \ln(k)/k$ the set $S(\mathbf{\Phi})$ is conjectured to be dominated by a single connected component w.h.p. That is, there is a subset $\mathcal{C} \subset S(\mathbf{\Phi})$ that induces a connected subgraph of the Hamming cube such that $|\mathcal{C}| = (1 - o(1))|S(\mathbf{\Phi})|$ w.h.p. In fact, Markov chains such as 'Glauber dynamics' are expected to mix rapidly on the dominant component $\mathcal{C}$, and the two statements in Hypothesis 1.3 are expected to be true for $t = 1$. However, none of these statements has been proven rigorously.

If in fact the set $S(\mathbf{\Phi})$ is so 'uniform' that simple Markov chains mix rapidly on $\mathcal{C}$, then it is plausible to expect that simple local search algorithms such as simulated annealing might find a satisfying assignment in polynomial time. Indeed, a specifically tailored local search algorithm is rigorously known to succeed up to $r \sim r_d$ w.h.p. [12]. This is, for general k, the largest density for which a polynomial time algorithm is known to find a satisfying assignment w.h.p.

For $r_d < r < r_s - O(1)$ a phenomenon called *dynamic replica symmetry breaking* ('dRSB') complicates matters. This phenomenon was (to an extent) established rigorously in [2]. The set $S(\mathbf{\Phi})$ does not have a dominant component anymore, but shatters into exponentially many, exponentially small components whose pairwise Hamming distance is as large as $\Omega(n)$ w.h.p. Moreover, each of the small components is 'rigid' in the sense that almost all variables take the same value in all satisfying assignments in that component. If we let $\mathcal{C}_1, \ldots, \mathcal{C}_N$ denote the components of $S(\mathbf{\Phi})$, then the Gibbs measure $g_{\mathbf{\Phi}}$ can be decomposed as

$$g_{\mathbf{\Phi}} = \sum_{i=1}^{N} c_i \cdot g_{\mathcal{C}_i},$$

where $c_i = |\mathcal{C}_i|/|S(\mathbf{\Phi})|$ is the relative size of component $\mathcal{C}_i$, and $g_{\mathcal{C}_i}$ is the unifom distribution over $\mathcal{C}_i$. (That is, $g_{\mathcal{C}_i}$ assigns probabiliy $1/|\mathcal{C}_i|$ to each $\sigma \in \mathcal{C}_i$.) In other words, one can think of $g_{\mathbf{\Phi}}$ as consisting of two 'layers': firstly, there is a distribution over the components $\mathcal{C}_1, \ldots, \mathcal{C}_N$ induced by the coefficients c_i, and secondly there is the uniform distribution $g_{\mathcal{C}_i}$ inside of each component. Of course, if the sizes $|\mathcal{C}_i|$ of the components were concentrated (on a logarithmic scale), then the Gibbs measure would be 'close' (say, within $o(1)$ in total variation distance) to the measure

---

[3] Actually the Gibbs measure comes with a temperature parameter, which we have omitted for the sake of simlicity. The uniform distribution on $S(\Phi)$ is the Gibbs measure at zero temperature.

$$u_{\mathbf{\Phi}} = \frac{1}{N}\sum_{i=1}^{N} g_{\mathcal{C}_i}$$

without the leading coefficients c_i w.h.p. The two statements of Hypothesis 1.3 are expected to be true for $t = 1$ for densities r such that the two measures $g_{\mathbf{\Phi}}$ and $u_{\mathbf{\Phi}}$ are close w.h.p.

The paper [20] claims a further transition called *condensation* or *one-step replica symmetry breaking* at $r \sim r_c = 2^k \ln 2 - \frac{3}{2}\ln 2$ where the measures $g_{\mathbf{\Phi}}$ and $u_{\mathbf{\Phi}}$ cease to be close w.h.p. That is, for $r_c < r < r_s$ the coefficients c_i are no longer concentrated, but there is a subset $\mathcal{D} \subset [N]$ of size $\mathcal{D} = \exp(o(n))$ such that $\sum_{i \in \mathcal{D}} c_i = 1 - o(1)$ w.h.p. For densities $r_c < r < r_s$ Hypothesis 1.3 is not expected to hold, even for $t = 1$. However, the existence of the condensation transition has not been established rigorously. In fact, r_c exceeds the density $2^k \ln 2 - k$ for which $\mathbf{\Phi}$ is known to be satisfiable [5]. According to [20] the threshold for the existence of satisfying assignments is $r_s = 2^k \ln 2 - \frac{1}{2}(1 + \ln 2)$.

This picture suggests that `BPdec` will not succeed to find satisfying assignments for densities $r_c < r < r_s$ [20]. However, Theorem 1.1 shows that `BPdec` fails w.h.p. for densities $r = O(2^k/k)$ even below $r_d \sim 2^k \ln(k)/k$. As pointed out (non-rigorously) in [27], the reason for this is that the above picture neglects the effect of the decimation process. Recall that in each iteration `BPdec` decimates the formula by assigning a truth value to one variable (for good) and simplifying. Hence, as `BPdec` proceeds the lenghts of the clauses decrease. In other words, decimation imposes stronger contraints on the remaining, unassigned variables. As a consequence, the decimation process has a similar effect as increasing the clause/variable density and indeed, according to [27], decimation brings the condensation phenomenon about even for densities $r = O(2^k/k) < r_c$. (As we will show in the full version of this work, the techniques of the present paper can be used to establish this assertion rigorously.) Thus, even though for densities $r < r_c$ Hypothesis 1.3 may be true for $t = 1$, it ceases to hold for $t \geq (1 - r/2^k)n$.

Finally, let us comment on Survey Propagation guided decimation, another message passing algorithm for finding satisfying assignments of $\mathbf{\Phi}$ [8, 21]. Recall that `BPdec` is an attempt at implementing the idealized Experiment 1.2. In the present notation, Experiment 1.2 samples a satisfying from the Gibbs measure $g_{\mathbf{\Phi}}$. Clearly, one could think of a similar experiment that yields a sample from the measure $u_{\mathbf{\Phi}}$ by 'just' replacing the marginals $M_{x_t}(\mathbf{\Phi}_{t-1})$ of $g_{\mathbf{\Phi}_{t-1}}$ by the corresponding marginals $U_{x_t}(\mathbf{\Phi}_{t-1})$ of $u_{\mathbf{\Phi}_{t-1}}$. Survey Propagation guided decimation is basically an attempt at implementing this latter experiment by means of a message passing algorithm.

The motivation behind Survey Propagation is that the statement about the marginals $U_{x_t}(\mathbf{\Phi}_{t-1})$ corresponding to Hypothesis 1.3 might be true for a larger range of r (and t). For, as shown in [20, 27], the basic obstacle to Hypothesis 1.3 being true is the condensation phenomenon, i.e., the fact that the measure $g_{\mathbf{\Phi}}$ is dominated by a sub-exponential number of components $\mathcal{C}_i$. But by construction this cannot happen in the case of $u_{\mathbf{\Phi}}$, unless, of course, the *total* number N of components is sub-exponential.

The above discussion describes the asymptotic picture for general clause lenghts k. Its qualitiative conclusions can be verified to an extent already for $k \geq 10$ [2]. However, generally, random k-SAT for 'small' k (say, $k \leq 6$) appears to behave somewhat differently (see the experimental study [7]). Technically an analysis for, say, $k \leq 6$ is quite challenging, not least because the relative error in the estimate of the k-SAT threshold obtained through the second moment method is larger than for general k [5].

A.2 Variations on `BPdec`

There are various different ways of implementing the basic idea of combining Belief Propagation with a decimation procedure. The algorithm `BPdec` from Section 1 is the simplest but also arguably the most natural variant. Other variants used in experiments [19, 20, 27] differ in

a. the number ω of iterations of the BP operator used in Step 2, and/or

b. the order in which the variables are assigned.

While for the sake of simplicity `BPdec` assumes that the number of iterations is given by some fixed function $\omega = \omega(n)$, it could make sense to use a 'dynamic' criterion such as the difference between the message vectors $\mu[\ell]$ and $\mu[\ell - 1]$ in some norm. For instance, one could run BP until it reaches an (approximate) fixed point. The present analysis extends easily to this variant. Amending the proof accordingly is straightforward, because the key statement Proposition 5.1 on the dyanmics of BP holds for any ℓ. That said, according to Hypothesis 1.3 (which inspired the use of Belief Propagation in the first place), the number ω of iterations should not matter, so long as $1 \ll \omega \ll \ln n$.

With respect to b., sometimes the variable x for which $|\frac{1}{2} - \mu_x(\mathbf{\Phi}_{t-1}, \omega)|$ is largest is chosen to be assigned at iteration t. The reason given for this choice is essentially numerical stability: numerical errors might disperse 'hard' constraints such as unit clauses. However, conceptually the choice of the variable to be assigned next should be unimportant, because the (ideal) BP operator does respect such hard constraints implicitly.

An interesting open problem is to generalize the present analysis to Survey Propagation. Furthermore, it would be interesting to see if the techniques developed here to prove a positive result about `BPdec`.

B Proof of Proposition 3.2

We consider an additional variant of `BPdec` that receives the order π in which variables are to be decimated as an input parameter.

Algorithm B.1 `BPdec`(Φ, π)
Input: A k-SAT formula Φ on $V = \{x_1, \ldots, x_n\}$, a permutation $\pi \in S_n$.
Output: An assignment $\tau : V \to \{0, 1\}$.

```
1.  Let $\Phi_0 = \Phi$.
1.  For $t = 1, \ldots, n$ do
2.      Compute the BP results $\mu_x(\Phi_{t-1}, \omega)$.
3.      Let $\sigma(x_{\pi(t)}) = 1$ with probability $\mu_x(\Phi_{t-1}, \omega)$, and set $\sigma(x_{\pi(t)}) = 0$ otherwise.
4.      Obtain $\Phi_t$ from $\Phi_{t-1}$ by substituting the value $\sigma(x_{\pi(t)})$ for $x_{\pi(i)}$.
5.  Return the assignment $\sigma$.
```

Fix a k-CNF Φ that is $(\boldsymbol{\delta}, t, \xi)$-uniform for all $1 \leq t \leq T$. Let S_n be the set of all permutations on $[n]$. Let λ be the probability distribution on pairs $(\boldsymbol{\pi}, \boldsymbol{\sigma}) \in S_n \times \{0,1\}^V$ induced by choosing a permutation $\boldsymbol{\pi} \in S_n$ uniformly and letting $\boldsymbol{\sigma} = \texttt{BPdec}(\Phi, \boldsymbol{\pi})$. Then $\bar{\beta}_\Phi$ is the $\boldsymbol{\sigma}$-marginal of λ, whence

$$\bar{\beta}_\Phi(\mathcal{E}) = \lambda(S_n \times \mathcal{E}) \quad \text{for any } \mathcal{E} \subset \{0,1\}^V . \tag{16}$$

In order to study λ, we consider another distribution λ' on pairs $(\boldsymbol{\pi}, \boldsymbol{\sigma}') \in S_n \times \{0,1\}^V$ that is easier to analyze and that will turn out to be 'close' to λ. To define λ', let $\mathcal{B}_t$ be the set of all pairs (π, σ) such that (Φ, π, σ) is not (δ_t, t)-balanced. Moreover, let $\mathcal{B} = \bigcup_{t=0}^{T} \mathcal{B}_t$. The distribution λ' is induced by choosing a permutation $\boldsymbol{\pi}$ uniformly at random and running the following algorithm on $\Phi, \boldsymbol{\pi}$.

Algorithm B.2 `BPdec`$'(\Phi, \pi)$
Input: A k-SAT formula Φ on $V = \{x_1, \ldots, x_n\}$, a permutation $\pi \in S_n$.
Output: An assignment $\boldsymbol{\sigma}' : V \to \{0, 1\}$.

```
0.  Let $\Phi_0 = \Phi$.
1.  For $t = 1, \ldots, n$ do
2.      Compute the BP results $\mu_x(\Phi_{t-1}, \omega)$.
3.      If $(\Phi, \pi, \boldsymbol{\sigma}')$ is $(\delta_t, t)$-balanced, then
            let $\boldsymbol{\sigma}'(x_{\pi(t)}) = 1$ with probability $\mu_{x_{\pi(t)}}(\Phi_{t-1}, \omega))$, and set $\boldsymbol{\sigma}'(x_{\pi(t)}) = 0$ otherwise.
        else
            let $\boldsymbol{\sigma}'(x_{\pi(t)}) = \zeta$ with probability $\frac{1}{2}$ for $\zeta = 0, 1$.
4.      Obtain $\Phi_t$ from $\Phi_{t-1}$ by substituting the value $\boldsymbol{\sigma}'(x_{\pi(i)})$ for $x_{\pi(i)}$ and simplifying.
5.  Output the assignment $\boldsymbol{\sigma}'$.
```

Roughly speaking, `BPdec`$'$ disregards the outcome μ of BP if it strays too far from the 'flat' vector $\frac{1}{2}\mathbf{1}$. We claim that λ and λ' are related as follows. For $\mathcal{F} \subset S_n \times \{0,1\}^V$ we let $\mathcal{F}_T$ be the set of all pairs (π, σ) for which there is $(\pi^*, \sigma^*) \in \mathcal{F}$ such that $\pi_*(t) = \pi(t)$ and $\sigma^*(x_{\pi(t)}) = \sigma(x_{\pi(t)})$ for all $1 \leq t \leq T$. Intuitively, $\mathcal{F}_T$ is the projection of $\mathcal{F}$ on the first T steps. We are going to show that

$$\lambda(\mathcal{F}) \quad \leq \quad \lambda'(\mathcal{F}_T) + \lambda'(\mathcal{B}) \quad \text{for any } \mathcal{F} \subset S_n \times \{0,1\}^V . \tag{17}$$

To establish (17), note that by construction for any $(\pi,\sigma)\notin\mathcal{B}_t$ and any $\zeta\in\{0,1\}$ we have

$$\lambda\left[\boldsymbol{\sigma}(x_{\pi(t+1)})=\zeta|\boldsymbol{\pi}=\pi\wedge\forall s\leq t:\boldsymbol{\sigma}(x_{\pi(s)})=\sigma(x_{\pi(s)})\right]$$
$$=\quad\lambda'\left[\boldsymbol{\sigma}'(x_{\pi(t+1)})=\zeta|\boldsymbol{\pi}=\pi\wedge\forall s\leq t:\boldsymbol{\sigma}'(x_{\pi(s)})=\sigma(x_{\pi(s)})\right].$$

Consequently, for any pair $(\pi,\sigma)\notin\mathcal{B}$ we get

$$\lambda\left[\forall t\leq T:\boldsymbol{\pi}(t)=\pi(t)\wedge\boldsymbol{\sigma}(t)=\sigma(t)\right]=\lambda'\left[\forall t\leq T:\boldsymbol{\pi}(t)=\pi(t)\wedge\boldsymbol{\sigma}'(t)=\sigma(t)\right].\tag{18}$$

In particular, $\lambda(\mathcal{B})=\lambda'(\mathcal{B})$. Hence, for any event $\mathcal{F}$ we obtain

$$\lambda(\mathcal{F})\quad\leq\quad\lambda(\mathcal{F}_T)\leq\lambda(\mathcal{F}_T\setminus\mathcal{B})+\lambda(\mathcal{B})\overset{(18)}{=}\lambda'(\mathcal{F}_T\setminus\mathcal{B})+\lambda'(\mathcal{B})\leq\lambda'(\mathcal{F}_T)+\lambda'(\mathcal{B}),$$

thereby proving (17).

Let λ'' be the uniform probability distribution on $S_n\times\{0,1\}^V$, and let $(\boldsymbol{\pi},\boldsymbol{u})$ be a uniformly random pair. We are going to relate λ' and λ''. To this end, let A_t be equal to one if $(\pi,\sigma)\notin\mathcal{B}_t$ and $x_{\pi(t)}$ is (δ_t,t)-biased in (Φ,π,σ). In addition, let $A=\sum_{t\leq T}A_t$. We claim that for any pair (π,σ)

$$\lambda'\left[\forall t\leq T:\boldsymbol{\pi}(t)=\pi(t)\wedge\boldsymbol{\sigma}'(x_{\pi(t)})=\sigma(x_{\pi(t)})\right]$$
$$\leq\quad\lambda''\left[\forall t\leq T:\boldsymbol{\pi}(t)=\pi(t)\wedge\boldsymbol{u}(x_{\pi(t)})=\sigma(x_{\pi(t)})\right]\cdot 2^{A(\pi,\sigma)}\prod_{t\leq T}1+2\delta_t.\tag{19}$$

To see this, fix any pair (π,σ) and let $\mathcal{L}_t=\left\{(\pi^*,\sigma^*):\pi^*(t)=\pi(t)\wedge\sigma^*(x_{\pi(t)})=\sigma(x_{\pi(t)})\right\}$. Let $1\leq t\leq T$. We can bound the conditional probability $\lambda'\left[\mathcal{L}_t|\boldsymbol{\pi}(t)=\pi(t)\wedge\bigwedge_{s<t}\mathcal{L}_s\right]$ as follows.

Case 1: $(\pi,\sigma)\in\mathcal{B}_t$**.** In this case (Φ,π,σ) is not (δ_t,t)-balanced. Therefore, step 3 of `BPdec'` chooses the value $\boldsymbol{\sigma}'(x_{\pi(t)})$ uniformly. Hence, the event $\boldsymbol{\sigma}'(x_{\pi(t)})=\sigma(x_{\pi(t)})$ occurs with probability $\frac{1}{2}$.

Case 2: $(\pi,\sigma)\notin\mathcal{B}_t$ **and** $A_t(\pi,\sigma)=0$**.** Since (Φ,π,σ) is (δ_t,t)-balanced, step 3 of `BPdec'` uses the BP marginals $\mu_{x_{\pi(t)}}(\zeta)$ in order to assign $x_{\pi(t)}$. Because $A_t(\pi,\sigma)=0$, the variable $x_{\pi(t)}$ is not (δ_t,t)-biased, whence $\mu_{x_{\pi(t)}}(\zeta)\leq\frac{1}{2}+\delta_t$ for $\zeta=0,1$. Hence, the probability that $\boldsymbol{\sigma}'(x_{\pi(t)})=\sigma(x_{\pi(t)})$ is bounded by $\frac{1}{2}+\delta_t$.

Case 3: $A_t(\pi,\sigma)=1$**.** In this case we just use the trivial fact that the probability of the event $\boldsymbol{\sigma}'(x_{\pi(t)})=\sigma(x_{\pi(t)})$ is bounded by $1\leq 2(\frac{1}{2}+\delta_t)$.

In any case, we obtain the bound $\lambda'\left[\mathcal{L}_t|\boldsymbol{\pi}(t)=\pi(t)\wedge\bigwedge_{s<t}\mathcal{L}_s\right]\leq 2^{A_t(\pi,\sigma)}(\frac{1}{2}+\delta_t)$. Consequently, as λ'' is the uniform distribution, we get

$$\frac{\lambda'\left[\mathcal{L}_t|\boldsymbol{\pi}(t)=\pi(t)\wedge\bigwedge_{s<t}\mathcal{L}_s\right]}{\lambda''\left[\mathcal{L}_t|\boldsymbol{\pi}(t)=\pi(t)\wedge\bigwedge_{s<t}\mathcal{L}_s\right]}\leq 2^{A_t(\pi,\sigma)}(1+2\delta_t).$$

Multiplying this estimate up for $t\leq T$, we obtain (19).

To put (19) to work, we need to estimate the random variable $A(\boldsymbol{\pi},\boldsymbol{\sigma}')$. Let $\Sigma_T=\sum_{s\leq T}\delta_s$. We claim that

$$\lambda'\left[A(\boldsymbol{\pi},\boldsymbol{\sigma}')>4(\Sigma_T+\xi n)\right]\leq\exp(-\xi n).\tag{20}$$

To see this, we are going to bound the probability that $A_t(\boldsymbol{\pi},\boldsymbol{\sigma}')=1$ given the values $\boldsymbol{\pi}(s)$, $\boldsymbol{\sigma}'(x_{\boldsymbol{\pi}(s)})$ for $1\leq s<t$.

Case 1: the event $\mathcal{B}_t$ occurs. Then $A_t=0$ by definition.

Case 2: the event $\mathcal{B}_t$ does not occur. In this case (Φ,π,σ) is (δ_t,t)-balanced, which means that no more than $\delta_t(n-t)$ variables are biased. Since the permutation $\boldsymbol{\pi}$ is chosen uniformly at random, the probability that $x_{\boldsymbol{\pi}(t)}$ is (δ_t,t)-biased is bounded by δ_t.

Thus, in either case the conditional probability of the event $A_t=1$ is bounded by δ_t. This implies that the random variable $A(\boldsymbol{\pi},\boldsymbol{\sigma}')=\sum_{t\leq T}A_t(\boldsymbol{\pi},\boldsymbol{\sigma}')$ is stochastically dominated by a sum of mutually independent Bernoulli variables with means $\delta_1,\ldots,\delta_T$. Therefore, (20) follows from Chernoff bounds.

Finally, combining (19) and (20), we see that

$$\begin{aligned}\lambda'\left[\mathcal{F}_T\right] &\leq \lambda'\left[A_T > 4(\Sigma_T + \xi n)\right] + \lambda'\left[\mathcal{F}_T \wedge A_T \leq 4(\Sigma_T + \xi n)\right] \\ &\leq \exp(-\xi n) + \lambda''\left[\mathcal{F}_T\right] \cdot 2^{4(\Sigma_T + \xi n)} \prod_{t \leq T} 1 + 2\delta_t \\ &\leq \lambda''\left[\mathcal{F}_T\right] \cdot \exp(6\Sigma_T + 4\xi n) + \exp(-\xi n) \qquad \text{for any } \mathcal{F} \subset S_n \times \{0,1\}^V. \qquad (21)\end{aligned}$$

Our assumption that Φ is (δ, t, ξ)-uniform ensures that $\lambda''(\mathcal{B}_t) \leq \exp(-10(\xi n + \Sigma_T))$ for all $t \leq T$. Together with (21), this implies that

$$\lambda'\left[\mathcal{B}_t\right] \leq \lambda''\left[\mathcal{B}_t\right] \exp(6\Sigma_T + 4\xi n) + \exp(-\xi n) \leq 2\exp(-\xi n) \qquad \text{for any } t \leq T.$$

Therefore,

$$\lambda'\left[\mathcal{B}\right] \leq 2T\exp(-\xi n) \leq \exp(-0.9\xi n). \qquad (22)$$

Putting things together, we obtain for any $\mathcal{E} \subset \{0,1\}^V$

$$\begin{aligned}\bar{\beta}_\Phi(\mathcal{E}) &= \lambda\left[S_n \times \mathcal{E}\right] \\ &\leq \lambda'\left[(S_n \times \mathcal{E})_T\right] + \lambda'\left[\mathcal{B}\right] \qquad \text{[due to (17)]} \\ &\leq \lambda'\left[(S_n \times \mathcal{E})_T\right] + \exp(-0.9\xi n) \qquad \text{[by (22)]} \\ &\leq \lambda''\left[(S_n \times \mathcal{E})_T\right] \exp(6(\Sigma_T + \xi n)) + \exp(-0.9\xi n) \qquad \text{[by (21)]} \\ &= \frac{|(S_n \times \mathcal{E})_T|}{n!2^n} \cdot \exp(6(\Sigma_T + \xi n)) + \exp(-0.9\xi n) \qquad \text{[as } \lambda'' \text{ is uniform]} \\ &\leq \frac{|\mathcal{E}|}{2^T} \cdot \exp(6(\Sigma_T + \xi n)) + \exp(-0.9\xi n) \qquad \text{[by the definition of } \mathcal{E}_T\text{]},\end{aligned}$$

as desired.

C Proof of Corollary 3.4

For $1 \leq t \leq T$ and a k-CNF Φ we let $X_t(\Phi)$ signify the number of pairs $(\pi, \sigma) \in S_n \times \{0,1\}^V$ such that (Φ, π, σ) is not (δ_t, t)-balanced. Then $X_t(\Phi) = 2^n n! \cdot \mathrm{P}\left[(\Phi, \boldsymbol{\pi}, \boldsymbol{u}) \text{ is } (\delta_t, t)\text{-balanced}\right]$, whence

$$\Phi \text{ is } (\delta, t, \xi)\text{-uniform iff } X_t(\Phi) \leq 2^n n! \cdot \exp(-\xi n - 4\Delta_t). \qquad (23)$$

With respect to the random formula $\boldsymbol{\Phi}$, Theorem 3.3 yields $\mathrm{E}\left[X_t(\boldsymbol{\Phi})\right] \leq 2^n n! \cdot \exp(-3\xi - 4\Delta_t)$, for a certain $\xi > 0$. Hence, Markov's inequality and the union bound entail

$$\mathrm{P}\left[\exists t \leq T : X_t(\boldsymbol{\Phi}) > \exp(-\xi - 4\Delta_t)\right] \leq n\exp(-2\xi n) \leq \exp(-\xi n),$$

and thus the assertion follows from (23).

D Proof of Theorem 1.1

Let us keep the symbols introduced in Theorem 3.3. Let $\mathcal{U}$ be the event that $\boldsymbol{\Phi}$ is $(\boldsymbol{\delta}, t, \xi)$-uniform for all $1 \leq t \leq T$. Let $\mathcal{S}$ be the event that $|S(\boldsymbol{\Phi})| \leq n \cdot \mathrm{E}|S(\boldsymbol{\Phi})|$, i.e., the number $|S(\boldsymbol{\Phi})|$ of satisfying assignments of $\boldsymbol{\Phi}$ does not exceed its expectation too much. By Corollary 3.4 and Markov's inequality, we have $\boldsymbol{\Phi} \in \mathcal{U} \cap \mathcal{S}$ w.h.p. Given that $\mathcal{U}$ and $\mathcal{S}$ occur, Proposition 3.2 implies

$$\bar{\beta}_{\boldsymbol{\Phi}}(S(\boldsymbol{\Phi})) \leq n \cdot \mathrm{E}|S(\boldsymbol{\Phi})| \cdot 2^{-T} \exp\left[4\sum_{t=1}^{T} \delta_t\right] + \exp(-\xi n/2). \qquad (24)$$

Simple calculations show $\sum_{t=1}^{T} \delta_t \sim n/(ck\exp(ckr/2^k))$ and $\mathrm{E}|S(\boldsymbol{\Phi})| = 2^n(1-2^{-k})^m \leq 2^n \exp(-rn/2^k)$. Plugging these estimates into (24) and recalling that $T = (1 - r/2^k)n$, we get

$$\begin{aligned}\bar{\beta}_{\boldsymbol{\Phi}}(S(\boldsymbol{\Phi})) &\leq \exp\left[(1+o(1))n\left(\ln 2 - \frac{r}{2^k} + \frac{\exp(-ckr/2^k)}{ck}\right) - T\right] + \exp(-\xi n/2) \\ &\leq \exp\left[(1+o(1))n\left(\ln 2 - 1 + \frac{\exp(-ckr/2^k)}{ck}\right) - T\right] + \exp(-\xi n/2) \leq \exp(-\Omega(n)),\end{aligned}$$

provided that $kr/2^k \geq \rho_0$ for a large enough constant $\rho_0 > 0$. Finally, Theorem 1.1 follows from Fact 3.1.

E Proof of Proposition 4.2

As in Proposition 4.2 we let $\delta_s = \exp(-c(1-s/n)k)$ for a small enough $c > 0$ and any $1 \leq s \leq n$, and we assume that $\rho = kr/2^k \geq \rho_0$ for some large enough number ρ_0 satisfying $\rho_0 \geq \exp(1/c)$. Let $T = (1-\rho/k)n$ and suppose that $1 \leq t \leq T$. Set $\theta = 1 - t/n$ and observe that $\theta k \geq \rho$. Let $\delta = \delta_t$. Let $\mathbf{\Phi}_t$ be the formula obtain from $\mathbf{\Phi}$ by substituting the value 'true' for $x_1, \ldots, x_{t-1}$ and simplifying. Since the δ_s form a geometric sequence, we have

$$\Sigma_t \quad = \quad \sum_{s \leq t} \delta_s \sim \frac{n}{ck\exp(c\theta k)}.$$

Observe that

$$\theta\delta n \quad > \quad 10^{15}\Sigma_t$$

if $\rho \geq \rho_0$ is chosen sufficiently large.

Let $k_1 = \sqrt{c}\theta k$ (with $c > 0$ as in (6)). Moreover, for a variable $x \in V_t$ and a set $Q \subset V_t$ let

$$\begin{aligned} N_1(x,Q) &= \{b \in N(x) : |N(b) \setminus Q| \geq 0.01(1-t/n)k, |N(b) \cap Q \setminus x| = 1\}, \\ N_{>1}(x,Q) &= \{b \in N(x) : |N(b) \setminus Q| \geq 0.01(1-t/n)k, |N(b) \cap Q \setminus x| > 1\} \end{aligned}$$

To prove Proposition 4.2, we will study two slightly different models of random k-CNFs. In the first model $\mathbf{\Phi}'$, we obtain a k-CNF by including each of the $(2n)^k$ possible clauses over $V = \{x_1, \ldots, x_n\}$ with probability $p = m/(2n)^k$ independently. Thus, $\mathbf{\Phi}'$ is a random set of clauses, and $\mathrm{E}|\mathbf{\Phi}'| = m$. In the second model, we choose a *sequence* $\mathbf{\Phi}''$ of m independent k-clauses $\mathbf{\Phi}''(1), \ldots, \mathbf{\Phi}''(m)$. Thus, the probability of each individual sequence is $(2n)^{-km}$. The sequence $\mathbf{\Phi}''$ corresponds to the k-CNF $\{\mathbf{\Phi}''(1), \ldots, \mathbf{\Phi}''(m)\}$ with *at most* m clauses. (Recall that we defined a k-CNF as a *set* of clauses.) The following well-known fact relates $\mathbf{\Phi}$ to $\mathbf{\Phi}'$, $\mathbf{\Phi}''$.

Fact E.1 *Let $\mathcal{E}$ be any property of random k-CNFs. Then*

$$\begin{aligned} \mathrm{P}\left[\mathbf{\Phi} \in \mathcal{E}\right] &\leq C\sqrt{m} \cdot \mathrm{P}\left[\mathbf{\Phi}' \in \mathcal{E}\right], \\ \mathrm{P}\left[\mathbf{\Phi} \in \mathcal{E}\right] &\leq (1+o(1)) \cdot \mathrm{P}\left[\mathbf{\Phi}'' \in \mathcal{E}\right]. \end{aligned}$$

Proof. The total number of clauses in $\mathbf{\Phi}'$ is binomially distributed with mean m. Stirling's formula implies that the probability that $\mathbf{\Phi}'$ has *precisely* m clauses is at least $\Omega(1/\sqrt{m})$. Furthermore, the expected number of pairs $1 \leq i < j \leq m$ such that $\mathbf{\Phi}''(i) = \mathbf{\Phi}''(j)$ is $o(1)$. Therefore, $\mathrm{P}\left[|\{\mathbf{\Phi}''(1), \ldots, \mathbf{\Phi}''(m)\}| = m\right] = 1 - o(1)$, and given that this occurs, the set $\{\mathbf{\Phi}''(1), \ldots, \mathbf{\Phi}''(m)\}$ has the same distribution as $\mathbf{\Phi}$. □

Due to Fact E.1, it suffices to prove that the statements **Q0–Q4** hold in either models $\mathbf{\Phi}'$, $\mathbf{\Phi}''$ with probability at least $1 - \exp(-11\Sigma_T)$.

E.1 Proof of Q0.

A direct computation shows that the expected number of redundant clauses is $O(1)$. Furthermore, modifying a single clause of $\mathbf{\Phi}''$ can change the total number of redundant clauses by at most two, i.e., the number of redundant clauses is 2-Lipschitz. Therefore, Azuma's inequality implies the desired bound. With respect to the degrees of the variables, simply note that the total number of clauses m equals the sum of the degrees of the variables divided by k.

E.2 Proof of Q1.

Working with the model $\mathbf{\Phi}'$, we are going to prove the bound on the number of variables in very short/very long clauses first. For any j the expected number of clauses of length j is

$$\lambda_j = 2^j \binom{k}{j} \theta^j (1-\theta)^{k-j} \cdot \frac{m}{2^k} = \frac{2^j \rho\theta n}{j} \binom{k-1}{j-1} \theta^{j-1} (1-\theta)^{k-j}.$$

Let's begin with short clauses of lengths $j < \theta k/10$. The total expected number of such clauses is bounded by

$$\begin{aligned}\sum_{j<\theta k/10} \lambda_j &\leq 2^{0.1\theta k}\rho\theta n \cdot \mathrm{P}\left[\mathrm{Bin}(k-1,\theta) < \theta k/10\right] \\ &\leq 2^{0.1\theta k}\rho\theta n \cdot \exp(-\theta k/3) \leq \theta\exp(-\theta k/4)n,\end{aligned}$$

provided that $\theta k \geq \rho$ is sufficiently large. Hence, by Chernoff bounds the probability that in total there are more than $\theta n \cdot 10^{-5}\delta/(\theta k)$ clauses of length $\leq 0.1\theta k$ is bounded by

$$\exp\left[-\frac{10^{-5}\delta}{2\theta k}\cdot\ln\left(\frac{10^{-5}\delta/(\theta k)}{\theta\exp(-\theta k/4)n}\right)\cdot\theta n\right] \quad\leq\quad \exp\left(-\frac{\delta}{3\cdot 10^5\theta k}\cdot\theta k\cdot\theta n\right) \leq \exp\left(-10^{-6}\theta\delta n\right).$$

Moreover, if there are fewer than $\theta n \cdot 10^{-5}\delta/(\theta k)$ short clauses, then these do not contain more than $\theta n \cdot 10^{-6}\delta$ variables.

To bound the number of variables in clauses of length $\geq 10\theta k$, we use the following argument. For a given $\mu > 0$ let $\mathcal{L}_\mu$ be the event that $\mathbf{\Phi}_t$ has μ clauses so that the sum of the lengths of these clauses is at least $\lambda = 10\theta k\mu$. Then

$$\mathrm{P}\left[\mathcal{L}_\mu\right] \leq \binom{m}{\mu}\binom{k\mu}{\lambda}\theta^\lambda\left(\frac{1}{2}+\theta\right)^{k\mu-\lambda}.$$

(Explanation: choose the μ clauses; out of the $k\mu$ literals that they contain, choose λ whose underlying variables have indices $> (1-\theta)n$; since none of the μ has been eliminated from $\mathbf{\Phi}_t$, the remaining $k\mu - \lambda$ literals must either be negative or their underlying variables are in V_t.) Thus,

$$\begin{aligned}\mathrm{P}\left[\mathcal{L}_\mu\right] &\leq \binom{m}{\mu}\binom{k\mu}{\lambda}(2\theta)^\lambda\left(\frac{1}{2}+\theta\right)^{k\mu} \\ &\leq \binom{m}{\mu}\left[\left(\frac{1}{2}+\theta\right)\left(\frac{\mathrm{e}}{5}\right)^{10\theta}\right]^{k\mu} \leq \left(\frac{\mathrm{e}n\rho}{k\mu}\right)^\mu\left[(1+2\theta)\left(\frac{\mathrm{e}}{5}\right)^{10\theta}\right]^{k\mu} \\ &\leq \left(\frac{\mathrm{e}n\rho}{k\mu}\right)^\mu\left[\exp(2\theta)\left(\frac{\mathrm{e}}{5}\right)^{10\theta}\right]^{k\mu} \leq \left(\frac{\mathrm{e}n\rho}{k\mu}\right)^\mu\left(\frac{\mathrm{e}}{4}\right)^{10\mu\theta k} \\ &\leq \left[\frac{\mathrm{e}n\rho}{k\mu}\left(\frac{\mathrm{e}}{4}\right)^{10\theta k}\right]^\mu \\ &= \left[\frac{\mathrm{e}n\rho\cdot 10\theta k}{k\lambda}\left(\frac{\mathrm{e}}{4}\right)^{10\theta k}\right]^\mu = \left[\frac{10\mathrm{e}\theta n\rho}{\lambda}\left(\frac{\mathrm{e}}{4}\right)^{10\theta k}\right]^{\lambda/(10\theta k)} \\ &= \left[\left(\frac{10\mathrm{e}\theta n\rho}{\lambda}\right)^{1/(10\theta k)}\left(\frac{\mathrm{e}}{4}\right)\right]^\lambda.\end{aligned}$$

Hence, for $\lambda > 10^{-6}\theta\delta n$ we get

$$\begin{aligned}\mathrm{P}\left[\mathcal{L}_\mu\right] &\leq \binom{m}{\mu}\binom{k\mu}{\lambda}(2\theta)^\lambda\left(\frac{1}{2}+\theta\right)^{k\mu} \\ &\leq \left[\left(\frac{1000\mathrm{e}\rho}{\delta}\right)^{1/10\theta k}\left(\frac{\mathrm{e}}{4}\right)\right]^\lambda \leq \left(\frac{\mathrm{e}}{3}\right)^\lambda \leq \exp(-10^{-10}\theta\delta n),\end{aligned}$$

provided that $\rho \geq \rho_0$ is sufficiently large. In summary, we have shown that

$$\begin{aligned}\mathrm{P}\left[\left|\left\{x \in V_t : \exists b \in N(x) : \frac{|N(b)|}{\theta k} \not\in [0.1, 10]\right\}\right| > 10^{-5}\theta\delta n\right] &\leq \exp(-10^{-11}\theta\delta n) \\ &\leq \exp(-100\Sigma_t). \qquad (25)\end{aligned}$$

Finally, we need to bound the number of variables $x \in V_t$ for which $\delta(\theta k)^3\sum_{b\in N(x)}2^{-|N(b)|} > 1$. Let

$$\mu_j = \rho 2^j\binom{k-1}{j-1}\theta^{j-1}(1-\theta)^{k-j+1} \qquad (26)$$

be the expected number of clauses of length j that contain a given variable $x \in V_t$ in the formula $\mathbf{\Phi}_t$. Let $X_j(x)$ be the actual number, and let $X_{jl}(x)$ be the number of such clauses where x appears in position l in that clause ($1 \leq l \leq j$). Then by Chernoff bounds

$$\mathrm{P}\left[X_{jl}(x) > 10(\mu_j + \delta^{-1}(\theta k)^{-5})\right] \quad \leq \quad \exp(-10/(\delta(\theta k)^5)).$$

Let $\zeta = \exp(-10/(\delta(\theta k)^5))$. For j, l let V_{jl} be the set of all variables $x \in V_t$ such that $X_{jl}(x) > 10(\mu_j + \delta^{-1}(\theta k)^{-5})$. Since the random variables $(X_{jl}(x))_{x\in V_t}$ are mutually independent, Chernoff bounds yield

$$\begin{aligned}\mathrm{P}\left[|V_{jl}| > \frac{\theta\delta n}{(\theta k)^9}\right] &\leq \exp\left[-\frac{\theta\delta n}{2(\theta k)^9} \cdot \ln\left(\frac{\theta\delta n}{(\theta k)^9 \zeta}\right)\right] \\ &\leq \exp\left[\frac{\theta\delta n}{3(\theta k)^9} \cdot \ln \zeta\right] \leq \exp\left[-\frac{\theta n}{(\theta k)^{15}}\right], \end{aligned} \tag{27}$$

provided that $\rho \geq \rho_0$ is sufficiently large. Furthermore, if $x \notin V_{jl}$ for any $1 \leq j \leq 10\theta k$ and any $1 \leq l \leq j$, then for $\rho \geq \rho_0$ sufficiently large

$$\begin{aligned}\sum_{b\in N(x):|N(b)|\leq 10\theta k} 2^{-|N(b)|} &\leq 10 \sum_{j\leq 10\theta k} 2^{-j}(\mu_j + \delta^{-1}(\theta k)^{-5}) \\ &\leq 10\delta^{-1}(\theta k)^{-5} + 10 \sum_{j\leq 10\theta k} 2^{-j}\mu_j \\ &\leq 10\delta^{-1}(\theta k)^{-5} + 10\rho \leq 20\delta^{-1}(\theta k)^{-5} \leq \delta^{-1}(\theta k)^{-4}.\end{aligned}$$

Hence, the assertion follows from (27) and our bound on the number of variables in clauses of length $> 10\theta k$.

E.3 Proof of Q2.

We will work with the model $\mathbf{\Phi}'$ once more. Let $Q \subset V_t$ be a set of size $|Q| \leq \theta\delta n$. We first deal with the sums over $b \in N_{>1}(x, Q)$ for $x \in V_t$. Given our bound on the total number of variables in clauses of lengths smaller than $0.1\theta k$ or greater than $10\theta k$, we may assume that x occurs in no such clause. For a variable x we let $\mathcal{Q}(x, i, j, l)$ be the number of clauses b such that clause b contains x (or $\bar{x}$) in position i in the decimated formula, $0.1\theta k \leq |N(b)| = j \leq 10\theta k$, and $|N(b) \cap Q \setminus x| = l$. Then $\mathcal{Q}(x, i, j, l)$ is binomial and for $l \geq 2$ we have for $\rho \geq \rho_0$ sufficiently large

$$\mathrm{E}\mathcal{Q}(x, i, j, l) \leq \binom{j}{l}\delta^l \mu_j \leq \rho\binom{j}{l}\delta^l 2^j \leq 2^{j-l}\delta^{1.9},$$

with μ_j as defined in (26). Hence, by Chernoff for $j - l > k_1$ we get

$$\mathrm{P}\left[\mathcal{Q}(x, i, j, l) > 10 \cdot 2^{j-l}\delta^{1.9}\right] \leq \exp(-2^{j-l}\delta^{1.9}) \leq \exp(-2^{k_1}\delta^{1.9}) \leq \exp(-\exp(c^{2/3}\theta k)).$$

Let $\mathcal{Z}(i, j, l)$ be the number of variables x for which $\mathcal{Q}(x, i, j, l) > 10 \cdot 2^{j-l}\delta^{1.9}$. Then $\mathcal{Z}(i, j, l)$ is binomial (because we fix the position i where x occurs) and $\mathrm{E}\mathcal{Z}(i, j, l) \leq \theta n \exp(-\exp(c^{2/3}\theta k))$. Hence, by Chernoff bounds

$$\begin{aligned}\mathrm{P}\left[\mathcal{Z}(i, j, l) > \theta\delta n/(\theta k)^4\right] &\leq \exp\left[-\frac{\theta\delta n}{(\theta k)^4}\ln\left(\frac{\theta\delta n/(\theta k)^4}{\theta n\exp(-\exp(c^{2/3}\theta k))}\right)\right] \\ &\leq \exp\left[-\frac{\theta\delta n}{(\theta k)^4} \cdot \exp(c^{2/3}\theta k)\right].\end{aligned} \tag{28}$$

Furthermore, if x is a variable such that $\mathcal{Q}(x, i, j, l) \leq 10 \cdot 2^{j-l}\delta^{1.9}$ for all $1 \leq i \leq j$, $2 \leq l \leq j - k_1$, and $0.1\theta k \leq j \leq 10\theta k$, then for $\rho \geq \rho_0$ sufficiently large

$$\sum_{b\in N_{>1}(x,Q)} 2^{|N(b)\cap Q\setminus x|-|N(b)|} \quad \leq \quad \sum_{0.1\theta k\leq j\leq 10\theta k}\ \sum_{2\leq l\leq j-k_1} 10j\delta^{1.9} \leq \delta^{1.8}.$$

Finally, by the union bound and (28) the probability that *there is* a set Q such that $\mathcal{Z}(i,j,l) > \theta\delta n/(\theta k)^4$ for any i,j,l is bounded by

$$\begin{aligned} n^2\binom{n}{\theta\delta n}\exp\left[-\frac{\theta\delta n}{(\theta k)^4}\cdot\exp(c^{2/3}\theta k)\right] &\leq \exp\left[o(n)+\theta\delta n(1-\ln(\theta\delta))-\frac{\theta\delta n}{(\theta k)^4}\cdot\exp(c^{2/3}\theta k)\right] \\ &\leq \exp\left[\theta\delta n\left[o(1)-2\ln\delta-\exp(c^{3/4}\theta k)\right]\right] \\ &\leq \exp\left[-100\Sigma_t\right], \end{aligned}$$

provided that $\rho \geq \rho_0$ is sufficiently large. This implies the claim regarding the middle sum in **Q2**.

The second part of the proof is about clauses $b \in N(x)$ such that $|N(b) \cap Q \setminus x| = 1$. In this case a similar argument applies, but we need to work with $l = 1$. We have

$$\mathrm{E}\mathcal{Q}(x,i,j,1) \leq j\delta\mu_j \leq \rho j\delta 2^j.$$

For $0.1\theta k \leq j \leq 10\theta k$ Chernoff bounds yield

$$\mathrm{P}\left[\mathcal{Q}(x,i,j,1) > 10\rho j\delta 2^j\right] \leq \exp(-\exp(c^{2/3}\theta k)).$$

Let $\mathcal{Z}(i,j)$ be the number of variables $x \in V_t$ such that $\mathcal{Q}(x,i,j,1) > 10\rho j\delta 2^j$. Then $\mathcal{Z}(i,j)$ is binomial with $\mathrm{E}\mathcal{Z}(i,j) \leq \theta n\exp(-\exp(c^{2/3}\theta k))$, and thus Chernoff bounds yield

$$\mathrm{P}\left[\mathcal{Z}(i,j) > \theta\delta n/(\theta k)^4\right] \quad\leq\quad \exp\left[-\theta\delta n\right].$$

Hence, taking the union bound over Q, i, j, we see that with probability $\geq 1-\exp(-100\Sigma_t)$ there are no more than $\theta\delta n/(\theta k)$ variables x such that $\mathcal{Q}(x,i,j,1) > 10\rho j\delta 2^j$ for any i,j. Moreover, if x is such that $\mathcal{Q}(x,i,j,1) > 10\rho j\delta 2^j$ for all i,j, then

$$\sum_{b\in N_1(x,Q)} 2^{-|N(b)|} \quad\leq\quad \sum_{0.1\theta k\leq j\leq 10\alpha k} 10j^2\rho\delta \leq (\theta k)^5\delta,$$

as desired.

We come to the third sum in **Q2**. Once more, we fix a set Q of size $|Q| \leq \theta\delta n$. As before, we may disregard variables that occur in clauses of lengths smaller than $0.1\theta k$ or greater than $10\theta k$. Thus, fix a clause length $0.1\theta k \leq j \leq 10\theta k$ and a position $1 \leq l \leq j$. For a variable $x \in V_t$ let $\mathcal{N}_+(x,j,l)$ be the number of clauses $b \in N(x)$ such that $|N(b)| = j$, x_i occurs in position l in b positively, and $|N(b) \cap Q \setminus x| \leq 1$. Define $\mathcal{N}_-(x,j,l)$ similarly for negative occurrences of x. Then $\mathrm{E}(\mathcal{N}_+ + \mathcal{N}_-) \leq \mu_j \leq 2^j\rho$ (with μ_j as in (26)). Furthermore, $\mathcal{N}_+(x,j,l), \mathcal{N}_-(x,j,l)$ are binomially distributed with identical means, because in $\mathbf{\Phi}$ each literal is positive/negative with probability $\frac{1}{2}$. Hence, for $j \geq 0.1\theta k$ Chernoff bounds yield

$$\begin{aligned} \mathrm{P}\left[|\mathcal{N}_+(x,j,l)-\mathcal{N}_-(x,j,l)| > 2^j\delta(\theta k)^{-3}\right] &\leq \exp\left[-\frac{(2^j\delta(\theta k)^{-3})^2}{3(2^j\delta(\theta k)^{-3}+\mu_j)}\right] \\ &\leq \exp\left[-\frac{2^j\delta^2}{4(\theta k)^6\rho}\right] \leq \exp(-\exp(\theta k/100)), \end{aligned}$$

provided that $c > 0$ is chosen sufficiently small. For different variables $x \in V_t$ the random variables $\mathcal{N}_+(x,j,l) - \mathcal{N}_-(x,j,l)$ are independent (because we fixed the position l where x occurs). Hence, the number $\mathcal{B}(j,l)$ of variables for which $|\mathcal{N}_+(x,j,l)-\mathcal{N}_-(x,j,l)| > 2^j\delta(\theta k)^{-3}$ is binomially distributed with mean $\leq \theta n\exp(-\exp(0.01\theta k))$. Consequently, Chernoff bounds give

$$\begin{aligned} \mathrm{P}\left[\mathcal{B}(j,l) > \theta\delta n/(\theta k)^3\right] &\leq \exp\left[-\frac{1}{2}\theta\delta n/(\theta k)^3\ln\left(\frac{\theta\delta n/(\theta k)^3}{\exp(-\exp(\theta k))\theta n}\right)\right] \\ &\leq \exp\left[-\frac{\theta\delta n}{(\theta k)^3}\cdot\exp(\theta k/101)\right] \leq \exp\left[-\theta\delta n\exp(\theta k/102)\right] \end{aligned}$$

provided that $\rho \geq \rho_0$ is sufficiently large. By the union bound, the probability that *there is* a set Q and j,l such that $\mathcal{B}(j,l) > \theta\delta n/(\theta k)^3$ is bounded by

$$\begin{aligned} n^2\binom{n}{\theta\delta n}\exp\left[-\theta\delta n\exp(\theta k/102)\right] &\leq \exp\left[o(n)+\theta\delta n(1-\ln(\theta\delta))-\theta\delta n\exp(\theta k/102)\right] \\ &\leq \exp\left[o(n)-2\theta\delta\ln(\delta)n-\theta\delta\exp(\theta k/102)n\right] \\ &\leq \exp\left[-100\Sigma_t\right]. \end{aligned}$$

Finally, if $x \notin \mathcal{B}(j,l)$ for any j, l, then

$$\begin{aligned}\left|\sum_{b\in N_{\le 1}(x)} \frac{\mathrm{sign}(x,b)}{2^{|N(b)|}}\right| &= \left|\sum_{j,l} 2^{-j}(\mathcal{N}_+(x,j,l) - \mathcal{N}_-(x,j,l))\right| \\ &\le \sum_{j,l} 2^{-j}\left|\mathcal{N}_+(x,j,l) - \mathcal{N}_-(x,j,l)\right| \\ &\le \sum_{j,l} 2^{-j}\cdot 2^j\delta(\theta k)^{-3} \le 100\delta/(\theta k),\end{aligned}$$

provided that $\rho \ge \rho_0$ is large enough.

E.4 Proof of Q3.

We carry the proof out in the model $\mathbf{\Phi}''$. Let $0.01 \le z \le 1$ and let Q be a set of size $|Q| = \theta q n$ with $0.01\delta \le q \le 100\delta$. In the proof of **Q1** we actually showed that

$$\sum_{b:\frac{|N(b)|}{\theta k}\notin[0.1,10]} |N(b)| \le 10^{-5}\theta\delta n \le 10^{-3}|Q|.$$

Therefore, here we only need to consider clauses b of lengths $0.1\theta k \le |b| \le 10\theta k$.

For integers $S, Z > 0$ let $\mathcal{E}(Q,S,Z)$ be the event that in $\mathbf{\Phi}_t$ contains a set $\mathcal{Z}$ of Z clauses such that $\sum_{b\in\mathcal{Z}}|N(b)| = S$, $0.1\theta k \le |N(b)| \le 10\theta k$, and $|N(b)\cap Q| \ge z|N(b)|$ for all $b \in \mathcal{Z}$. We claim that for given Q, S, Z

$$\mathrm{P}\left[\mathcal{E}(Q,S,Z)\right] \le \binom{m}{Z}\binom{kZ}{S}\binom{S}{zS}2^{S-kZ}\theta^S(1-\theta)^{kZ-S}q^{zS}.$$

For there are m clauses in total, out of which we choose Z; then, out of the kZ literal occurrences of these Z clauses we choose S whose underlying variables lies in V_t, an event that occurs with probability $\theta = |V_t|/n$ independently for each. Furthermore, all $kZ - S$ literals whose variables are in $V \setminus V_t$ must be negative, because otherwise the corresponding clauses would have been satisfied by setting $x_1, \ldots, x_{t-1}$ to true. Finally, out of the S literal occurrences in V_t a total of at least zS has an underlying variable from Q, which occurs with probability $q = |Q|/(\theta n)$ independently.

Hence, we obtain

$$\begin{aligned}\mathrm{P}\left[\mathcal{E}(Q,S,Z)\right] &\le \binom{m}{Z}2^{-kZ}\left[2^{1/z}\cdot\frac{\mathrm{e}}{z}\cdot q\right]^{zS}\cdot\binom{kZ}{S}\theta^S(1-\theta)^{kZ-S} \\ &\le \binom{m}{Z}2^{-kZ}\left[2^{1/z}\cdot\frac{\mathrm{e}}{z}\cdot q\right]^{zS} \le \binom{m}{Z}2^{-kZ}(Cq)^{zS}\end{aligned}\tag{29}$$

for a certain absolute constant $C > 0$, because $z \ge 0.01$. Since all clause lengths are assumed to be between $0.1\theta k$ and $10\theta k$, we obtain $0.1S/(\theta k) \le Z \le 10S/(\theta k)$. Therefore,

$$\binom{m}{Z}2^{-kZ} \le \left(\frac{\mathrm{e}m}{2^kZ}\right)^Z \le \left(\frac{\mathrm{e}\rho n}{kZ}\right)^Z\left(\frac{10\mathrm{e}\theta k\cdot\rho n}{kS}\right)^{10S/(\theta k)} \le \left(\frac{10\mathrm{e}\rho}{1.01q}\right)^{10S/(\theta k)}.\tag{30}$$

Since $q \le 100\delta = 100\exp(-c\theta k)$ and $\theta k \ge \rho$, we have $1/q \ge 100\rho$ for $\rho \ge \rho_0$ sufficiently large. Hence, (30) yields

$$\binom{m}{Z}2^{-kZ} \le q^{-20S/(\theta k)}.\tag{31}$$

Plugging (31) into (29), we obtain for $\theta k \ge \rho \ge \rho_0$ large enough and $S \ge 1.01|Q|/z$

$$\mathrm{P}\left[\mathcal{E}(Q,S,Z)\right] \le q^{-20S/(\theta k)}\cdot(Cq)^{zS} \le q^{0.9999zS} \le q^{1.009|Q|}.\tag{32}$$

Let $\mathcal{E}(q,S,Z)$ denote the event that *there is* a set $Q \subset V_t$ of size $|Q| = q\theta n$ such that $\mathcal{E}(Q,S,Z)$ occurs. Then for $S \ge 1.01q\theta n/z$ the union bound and (32) yield

$$\mathrm{P}\left[\mathcal{E}(q,S,Z)\right] \le \binom{\theta n}{q\theta n}q^{1.009q\theta n} \le \exp\left[q\theta n(1-\ln q - 1.009\ln q))\right] \le \exp\left[-0.008q\theta n\ln q\right].\tag{33}$$

Finally, let $\mathcal{E}$ be the event that there is a set $Q \subset V_t$ of size $0.001\theta\delta n \le 100\theta\delta n$ and $S \ge 1.01|Q|/z$, $Z > 0$ such that $\mathcal{E}(q, S, Z)$ occurs. Then (33) yields for $\rho \ge \rho_0$ large enough

$$\mathrm{P}\left[\mathcal{E}\right] \quad \le \quad n^3 \cdot \exp\left[0.008 \cdot 0.0001\theta\delta n \ln(100\delta)\right] \le \exp(-100\Sigma_t),$$

as desired.

E.5 Proof of Q4.

We are going to work with the probability distribution $\mathbf{\Phi}''$ (sequence of m independent clauses). Let $\mathcal{M}$ be the set of all indices $l \in [m]$ such that the lth clause of $\mathbf{\Phi}''(l)$ is still present in the decimated formula at time t, i.e., $\mathbf{\Phi}''(l)$ does not contain any of the variables $x_1, \ldots, x_t$ positively. For each $l \in \mathcal{M}$ let $\mathcal{L}(l)$ be the number of literals in $\mathbf{\Phi}''(l)$ whose underlying variable is in V_t. We may assume without loss of generality that for any $l \in \mathcal{M}$ the $\mathcal{L}(l)$ *leftmost* literals $\mathbf{\Phi}''(l, i)$, $1 \le i \le \mathcal{L}(l)$, are the ones with an underlying variable from V_t.

Let $Q \subset V_t$. Analyzing the operator Λ_Q directly is a little awkward. Therefore, we will decompose Λ_Q into a sum of several operators that are easier to investigate. For any clause length $0.1\theta k \le L \le 10\theta k$, and $1 \le i < j \le L$, $l \in \mathcal{M}$, and any distinct $x, y \in V_t$ we define

$$m_{xy}(i,j,l,L) = \begin{cases} 1 & \text{if } \mathcal{L}(l) = L \text{ and } (\mathbf{\Phi}''(l,i) = x \wedge \mathbf{\Phi}''(l,j) = y) \vee (\mathbf{\Phi}''(l,i) = \bar{x} \wedge \mathbf{\Phi}''(l,j) = \bar{y}), \\ -1 & \text{if } \mathcal{L}(l) = L \text{ and } (\mathbf{\Phi}''(l,i) = x \wedge \mathbf{\Phi}''(l,j) = \bar{y}) \vee (\mathbf{\Phi}''(l,i) = \bar{x} \wedge \mathbf{\Phi}''(l,j) = y), \\ 0 & \text{otherwise,} \end{cases}$$

while we let $m_{xx}(i,j,l) = 0$. Moreover, for $x, y \in V_t$ we let $m_{xy}(i,j,L) = \sum_{l \in [m]} m_{xy}(i,j,l,L)$. For a variable $x \in V_t$ we let $\mathcal{N}(x, Q)$ be the set of all $l \in \mathcal{M}$ such that $0.1\theta k \le \mathcal{L}(l) \le 10\theta k$ and the clause $\mathbf{\Phi}''(l)$ contains at most one literal whose underlying variable is in $Q \setminus x$. We are going to analyze the $\|\cdot\|_\square$-norm of the operators

$$\Lambda_Q^{ijL} : \mathbf{R}^{V_t} \to \mathbf{R}^{V_t}, \quad \Gamma \mapsto \left[\sum_{b \in \mathcal{N}(Q,x)} \sum_{y \in N(b) \setminus x} 2^{-L} m_{xy}(i,j,L)\, \Gamma_y\right]_{x \in V}.$$

The following well-known fact about the norm $\|\cdot\|_\square$ of matrices with diagonal entries equal to zero simplifies matters a little. For $A \subset V_t$ we let $\mathbf{1}_A \in \{0,1\}^{V_t}$ denote the indicator of A.

Fact E.2 *We have* $\left\|\Lambda_Q^{ijL}\right\|_\square \le 24 \max_{A,B \subset V : A \cap B = \emptyset} \left|\left\langle \Lambda_Q^{ijL} \mathbf{1}_A, \mathbf{1}_B \right\rangle\right|$.

Lemma E.3 *Fix* i, j, L, Q. *With probability at least* $1 - \exp(-\theta n)$ *we have* $\left\|\Lambda_Q^{ijL}\right\|_\square \le \delta^5 \theta n$.

Proof. Fix two sets $A, B \subset V_t$. For each $l \in \mathcal{M}$ and any $x, y \in V_t$ the two $0/1$ random variables

$$\sum_{(x,y) \in A \times B} \max\left\{m_{xy}(i,j,l,L), 0\right\}, \qquad \sum_{(x,y) \in A \times B} \max\left\{-m_{xy}(i,j,l,L), 0\right\}$$

are identically distributed, because the clause $\mathbf{\Phi}''_l$ is chosen uniformly at random. In effect, since the clauses $(\mathbf{\Phi}''(l))_{l \in [m]}$ are mutually independent, each of the two random variables

$$\begin{aligned} \mu(A,B) &= \sum_{l \in \mathcal{M}} \sum_{(x,y) \in A \times B} \mathbf{1}_{l \in \mathcal{N}(x,Q)} \max\left\{m_{xy}(i,j,l), 0\right\}, \\ \nu(A,B) &= \sum_{l \in \mathcal{M}} \sum_{(x,y) \in A \times B} \mathbf{1}_{l \in \mathcal{N}(x,Q)} \max\left\{-m_{xy}(i,j,l), 0\right\} \end{aligned}$$

is a sum of independent Bernoulli variables, and the distributions of $\mu(A,B)$ and $\nu(A,B)$ coincide. As each of the clauses $\mathbf{\Phi}''(l)$ is chosen uniformly, for each $l \in [m]$ the probability that $l \in \mathcal{M}$ and $\mathcal{L}(l) = L$ equals $\binom{k}{L}\theta^L(1-\theta)^{k-L}2^{L-k}$. Therefore, we obtain the bound

$$\begin{aligned} \mathrm{E}(\mu(A,B)) &= \mathrm{E}(\nu(A,B)) \le m\binom{k}{L}\theta^L(1-\theta)^{k-L}2^{L-k} = \frac{2^k \rho n}{k} \cdot \binom{k}{L}\theta^L(1-\theta)^{k-L}2^{L-k} \\ &= \frac{2^L \rho\theta n}{L}\binom{k-1}{L-1}\theta^{L-1}(1-\theta)^{k-L} \le \frac{2^L \rho\theta n}{L}. \end{aligned} \tag{34}$$

Since both $\mu(A,B)$ and $\nu(A,B)$ are sums of independent Bernoulli variables, (34) entails that

$$\mathrm{Var}(\mu(A,B)) = \mathrm{Var}(\nu(A,B)) \leq \mathrm{E}(\mu(A,B)) = \mathrm{E}(\nu(A,B)) \leq \frac{2^L\rho\theta n}{L}.$$

Hence, Chernoff bounds yield that for a certain constant $c > 0$ we have

$$\begin{aligned} \mathrm{P}\left[|\mu(A,B) - \mathrm{E}(\mu(A,B))| > c\sqrt{2^L\rho/L}\cdot\theta n\right] &= \mathrm{P}\left[|\nu(A,B) - \mathrm{E}(\nu(A,B))| > c\sqrt{2^L\rho/L}\cdot\theta n\right] \\ &\leq 16^{-\theta n}. \end{aligned} \tag{35}$$

Since this is true for any fixed A, B, and because

$$\left\langle \Lambda_Q^{ijL}\mathbf{1}_B, \mathbf{1}_A \right\rangle = 2^{-L}(\mu(A,B) - \nu(A,B)),$$

the union bound and (35) imply the assertion. □

To complete the proof of **Q4**, let

$$\Lambda'_Q = \sum_{0.1\theta k \leq L \leq 10\theta k} \sum_{i=1}^{L} \sum_{j \neq i} \Lambda_Q^{ijL}.$$

Thus, both Λ_Q and Λ'_Q are $V_t \times V_t$-matrices. For $(x,y) \in V_t \times V_t$ corresponding the entries of Λ'_Q, Λ_Q differ only if either x or y occurs in a redundant clause. Consequently, by **Q0** $\left\|\Lambda'_Q - \Lambda_Q\right\|_{\square} = o(n)$. Therefore, **Q4** follows from Lemma E.3 and the union bound over s, t, L, and Q.

F Proof of Theorem 4.3

This appendix contains a complete proof of Theorem 4.3. It is self-contained in the sense that it is independent of the outline given in Section 5. *Throughout this section, we keep the notation from Section 4 and the assumptions of Theorem 4.3.* To unclutter the notation, we let $\delta = \delta_t$. Moreover, let ρ be such that $r = 2^k\rho/k$.

Let $k_1 = \sqrt{c}\theta k$ (with $c > 0$ as in (6)).

F.1 Outline

We are going to trace the BP operator when applied to the initial point $\mu_{x\to a}(0)[0] = \mu_{x\to a}(1)[0] = \frac{1}{2}$ for all $x \in V_t$, $a \in N(x)$. Let $\mu[\ell] = \mathrm{BP}^\ell(\mu[0]) \in M(\Phi)$ be the result of the first ℓ iterations of BP. Let

$$\Delta_{x\to a}[\ell] = \mu_{x\to a}(1)[\ell] - \frac{1}{2}.$$

We say that $x \in V_t$ is ***ℓ-biased*** if there is a clause $a \in N(x)$ such that $|\Delta_{x\to a}[\ell]| \geq 0.1\delta$. Clearly, no variable is 0-biased. Let $B[\ell]$ be the set of all ℓ-biased variables. To prove Theorem 4.3, the core task will be to bound the sizes $|B[\ell]|$.

We are going to construct a sequence of supersets $T[\ell] \supset B[\ell]$ whose sizes are easier to estimate. Actually we will construct sets of variables $T_1[\ell]$, $T_2[\ell]$ and sets of clauses $T_3[\ell]$ inductively and let $T[\ell] = T_1[\ell] \cup T_2[\ell] \cup N(T_3[\ell])$. For $\ell = 0$ we let $T_1[0] = T_3[0] = \emptyset$. Moreover, let $T_2[0]$ be the set of all variables x such that there is a clause $b \in N(x)$ that is either redundant, or $|N(b)| < 0.1\theta k$, or $|N(b)| > 10\theta k$, or that satisfy $\delta(\theta k)^3 \sum_{b\in N(x)} 2^{-|N(b)|} > 1$.

To define $T[\ell+1]$ inductively for $\ell \geq 0$, we need a bit of notation: for $x \in V$ and $a \in N(x)$ we let

$$\begin{aligned} N_{\leq 1}(x\to a)[\ell+1] &= N_{\leq 1}(x, T[\ell]) \setminus a, \qquad (36) \\ P_{\leq 1}(x\to a)[\ell+1] &= \prod_{b\in N_{\leq 1}(x\to a)[\ell+1]} \frac{\mu_{b\to x}(1)[\ell]}{\mu_{b\to x}(0)[\ell]}, \\ N_{>1}(x\to a)[\ell+1] &= \{b \in N(x) \setminus (a \cup N_{\leq 1}(x, T[\ell])) : \mu_{b\to x}(0)[\ell] \neq 0\}, \\ P_{>1}(x\to a)[\ell+1] &= \prod_{b\in N_{>1}(x\to a)[\ell+1]} \frac{\mu_{b\to x}(1)[\ell]}{\mu_{b\to x}(0)[\ell]}. \end{aligned}$$

It is direction from their definition (2) that the terms $\mu_{b\to x}(0)\,[\ell]$ are strictly positive for all $b \in N_{\le 1}(x \to a)$, so that $P_{\le 1}\,(x \to a)\,[\ell+1]$ is well-defined.

The idea behind these definitions is the following. Assume for a moment that $\mu_{b\to x}(0)\,[\ell] \neq 0$ for all $b \in N(x)$. Then in order to show that $\Delta_{x\to a}\,[\ell+1] = \mu_{x\to a}(1)\,[\ell+1] - \frac{1}{2}$ is close to zero it suffices to verify that the ratio

$$\frac{\mu_{x\to a}(1)\,[\ell+1]}{\mu_{x\to a}(0)\,[\ell+1]} = \prod_{b\in N(x)\setminus a} \frac{\mu_{b\to x}(1)\,[\ell]}{\mu_{b\to x}(0)\,[\ell]} = P_{\le 1}\,(x \to a)\,[\ell+1] \cdot P_{>1}\,(x \to a)\,[\ell+1] \tag{37}$$

is close to one, because $\mu_{x\to a}(0)\,[\ell+1] + \mu_{x\to a}(1)\,[\ell+1] = 1$ by construction. Moreover, (37) is close to one if both factors on the r.h.s. are.

The set $T_1\,[\ell+1]$ contains all variables for which $P_{\le 1}\,(x \to a)\,[\ell+1]$ fails to be close enough to one:

$$T_1\,[\ell+1] = \left\{ x \in V : \max_{a\in N(x)} |P_{\le 1}\,(x \to a)\,[\ell+1] - 1| > 0.01\delta \right\}.$$

To also deal with the second product $P^1_{x\to a}\,[\ell+1]$, we define additional sets $T_2\,[\ell+1]$, $T_3\,[\ell+1]$. To define $T_2\,[\ell+1]$, let us say that a variable x is $(\ell+1)$-***harmless*** if it enjoys the following four properties.

H1. We have $\delta(\theta k)^3 \sum_{b\in N(x)} 2^{-|N(b)|} \le 1$, and $0.1\theta k \le |N(b)| \le 10\theta k$ for all $b \in N(x)$.

H2. $\sum_{b\in N_1(x,T[\ell])} 2^{-|N(b)|} \le (\theta k)^5 \delta$ and $\sum_{b\in N_{>1}(x,T[\ell]))} 2^{|N(b)\cap T[\ell]\setminus x| - |N(b)|} \le \delta/(\theta k)$.

H3. There is at most one clause $b \in N(x)$ such that $|N(b) \setminus T\,[\ell]\,| \le k_1$.

H4. $\left|\sum_{b\in N_{\le 1}(x)} \mathrm{sign}(x,b) \cdot 2^{-|N(b)|}\right| \le 0.01\delta$.

Let $H\,[\ell+1]$ signify the set of all $(\ell+1)$-harmless variables. Further, let $T_2(\ell+1)$ be the set of all variables x that have at least one of the following properties.

T2a. There is a clause $b \in N(x)$ that is either redundant, or $|N(b)| < 0.1\theta k$, or $|N(b)| > 10\theta k$.

T2b. $\delta(\theta k)^3 \sum_{b\in N(x)} 2^{-|N(b)|} > 1$.

T2c. One of the following is true: either

$$\sum_{b\in N_1(x,T[\ell])} 2^{-|N(b)|} \quad > \quad (\theta k)^5 \delta, \text{ or } \sum_{b\in N_{>1}(x,T[\ell]))} 2^{|N(b)\cap T[\ell]\setminus x| - |N(b)|} > \delta/(\theta k).$$

T2d. x occurs in more than 100 clauses from $T_3\,[\ell]$.

T2e. x occurs in a clause b that contains fewer than $3|N(b)|/4$ variables from $H\,[\ell]$.

Proposition 4.2 (and items **Q0** and **Q1** from Definition 4.1) ensure that there are only a very few variables that satisfy **H1**, **T2a**, or **T2b** w.h.p. We always include these few into the set $T_2\,[\ell+1]$ of ‘exceptional’ variables. Moreover, intuitively **H2** and **T2c–T2e** capture variables x that are highly exposed to the ‘exceptional’ set $T\,[\ell]$ from the previous round. Furthermore, we let

$$T_3\,[\ell+1] = \{a \in \Phi_t : |N(a)| \ge 100k_1 \wedge |N(a) \setminus T\,[\ell]\,| \le k_1\} \setminus T_3\,[\ell]$$

contain all clauses that consist almost entirely of ‘exceptional’ variables from $T\,[\ell]$, but *without* including the clauses from the previous set $T_3\,[\ell]$. Finally,

$$T\,[\ell+1] = T_1\,[\ell+1] \cup T_2\,[\ell+1] \cup N(T_3\,[\ell+1]).$$

In Section F.2 we will verify that $T\,[\ell]$ does indeed contain the set $B\,[\ell]$ of biased variables.

Proposition F.1 *We have* $B\,[\ell] \subset T\,[\ell]$ *for all* $\ell \ge 0$.

Furthermore, in Section F.3 we will establish the following bound on the size of $T\,[\ell]$.

Proposition F.2 *We have* $|T\,[\ell]| \le 0.1\theta\delta n$ *for all* $\ell \ge 0$.

Finally, in Section F.6 we will derive Theorem 4.3 from Proposition F.1 and Proposition F.2.

F.2 Proof of Proposition F.1

Lemma F.3 *Let x be a variable and let $b \in N(x)$ be a clause. Let $t_b = |N(b) \cap B[\ell] \setminus x|$. Then*

$$|1-\mu_{b\to x}(\zeta)[\ell]| \leq 2^{2-|N(b)|+t_b}\exp(\delta|N(b)|) \quad \textit{for } \zeta = 0, 1.$$

Proof. We may assume that clause b is not a tautology, because otherwise $\mu_{b\to x}(0)[\ell] = \mu_{b\to x}(1)[\ell] = 1$. We consider two cases. The first case is that x appears negatively in b. Since for any $y \in N(b) \setminus x$ we have $\mu_{y\to b}(1)[\ell] = \frac{1}{2} + \Delta_{y\to b}[\ell]$ and $\mu_{y\to b}(0)[\ell] + \mu_{y\to b}(1)[\ell] = 1$, we see that

$$\mu_{y\to b}\left(\frac{1-\operatorname{sign}(y,b)}{2}\right)[\ell] = \frac{1}{2} - \operatorname{sign}(y,b)\,\Delta_{y\to b}[\ell].$$

Therefore,

$$\begin{aligned}
1-\mu_{b\to x}(1)[\ell] &= \prod_{y\in N(b)\setminus x} \frac{1}{2} - \operatorname{sign}(y,b)\,\Delta_{y\to b}[\ell] \\
&= 2^{1-|N(b)|}\prod_{y\in N(b)\setminus x} 1 - 2\operatorname{sign}(y,b)\,\Delta_{y\to b}[\ell] \\
&\leq 2^{1-|N(b)|}\cdot 2^{t_b}\cdot \prod_{y\in N(b)\setminus(x\cup B[\ell])} 1+2|\Delta_{y\to b}[\ell]| \leq 1+2^{1-|N(b)|+t_b}\cdot(1+\delta)^{|N(b)|} \\
&\leq 2^{1-|N(b)|+t_b}\cdot\exp(|N(b)|\delta).
\end{aligned}$$

Changing one sign in the above computation, we also see that

$$\mu_{b\to x}(1)[\ell] \quad\leq\quad 1+2^{2-|N(b)|+t_b}\cdot\exp(\delta|N(b)|).$$

In the second case where x appears positively in b the computation is analogous. □

Corollary F.4 *Let x be a variable and let $\mathcal{T} \subset N(x)$ be a set of clauses. For each $b \in \mathcal{T}$ let $t_b = |b \cap B[\ell] \setminus x|$. Assume that $t_b < |N(b)| - 2$ for all $b \in \mathcal{T}$. Then $\mu_{b\to x}(0)[\ell] > 0$ for all $b \in \mathcal{T}$ and*

$$\left|\ln\prod_{b\in\mathcal{T}}\frac{\mu_{b\to x}(1)[\ell]}{\mu_{b\to x}(0)[\ell]}\right| \leq \sum_{b\in\mathcal{T}} 2^{3-|N(b)|+t_b}\exp(\delta|N(b)|).$$

Proof. The fact that $\mu_{b\to x}(0)[\ell] > 0$ for all $b \in \mathcal{T}$ follows directly from the definition (2) of $\mu_{b\to x}$ and the assumption that $t_b < |N(b)| - 2$. Multiplying the terms from Lemma F.3 up for $b \in \mathcal{T}$, we see that

$$\prod_{b\in\mathcal{T}}\frac{\mu_{b\to x}(1)[\ell]}{\mu_{b\to i}(0)[\ell]} \quad\leq\quad \prod_{b\in\mathcal{T}} 1+2^{2-|N(b)|+t_b}\exp(\delta|N(b)|) \leq \exp\left[\sum_{b\in\mathcal{T}} 2^{2-|N(b)|+t_b}\exp(\delta|N(b)|)\right].$$

A similar computation yields the lower bound. □

Corollary F.5 *Suppose that $x \in H[\ell]$ is harmless and that $a \in N(x)$ is a clause such that $|N(a) \setminus T[\ell-1]| \leq k_1$. Then $|\Delta_{x\to a}[\ell]| \leq 0.01$.*

Proof. Condition **H3** ensures that $|N(b) \setminus T[\ell-1]| > k_1$ for any $b \in N(x) \setminus a$, and thus $\mu_{b\to x}(0)[\ell] > 0$. Moreover, **H2** yields

$$\sum_{b\in N_{>1}(x,T[\ell])} 2^{|N(b)\cap T[\ell-1]\setminus x|-|N(b)|} \quad\leq\quad \delta.$$

Since $\delta \leq \exp(-c\rho)$, Corollary F.4 applied to $\mathcal{T} = N_{>1}(x \to a, \ell) = N_{>1}(x, T[\ell]) \setminus a$ implies

$$|P_{>1}(x\to a)[\ell] - 1| \leq 10^{-4},$$

provided that $\rho \geq \rho_0$ is sufficiently large.

Furthermore, since by **H1** all clauses $b \in N(x)$ have lengths $0.1\theta k \leq |N(b)| \leq 10\theta k$, we have $N_{\leq 1}(x, T\,[\ell-1]) \supset N_1(x, T\,[\ell-1])$. Therefore,

$$|P_{\leq 1}\,(x \to a)\,[\ell] - 1| \quad = \quad \prod_{b \in N_{\leq 1}(x,T[\ell-1]) \setminus N_1(x,T[\ell-1])} \frac{\mu_{b\to x}(1)\,[\ell-1]}{\mu_{b\to x}(0)\,[\ell-1]} \cdot \prod_{b \in N_1(x,T[\ell-1])} \frac{\mu_{b\to x}(1)\,[\ell-1]}{\mu_{b\to x}(0)\,[\ell-1]}.$$

With respect to the second product, **H2** and Lemma F.3 yield for $\rho \geq \rho_0$ large

$$\begin{aligned} \left| \ln \prod_{b \in N_1(x,T[\ell-1])} \frac{\mu_{b\to x}(1)\,[\ell-1]}{\mu_{b\to x}(0)\,[\ell-1]} \right| &\leq 8 \sum_{b \in N_1(x,T[\ell-1])} 2^{-|N(b)|} \exp(\delta |N(b)|) \\ &\leq 16(\theta k)^5 \delta \leq 10^{-6}, \end{aligned}$$

because $|N(b)| \leq 10\theta k$ for all b and $\delta = \exp(-c\theta k)$. Furthermore, letting $\mathcal{N} = N_{\leq 1}(x, T\,[\ell-1]) \setminus N_1(x, T\,[\ell-1])$, we obtain from **H1** and **H4**

$$\begin{aligned} \left| \ln \prod_{b \in \mathcal{N}} \frac{\mu_{b\to x}(1)\,[\ell-1]}{\mu_{b\to x}(0)\,[\ell-1]} \right| &\leq \left| \sum_{b \in \mathcal{N}} \mathrm{sign}(x,b) 2^{1-|N(b)|} \right| \cdot \exp(10\theta k \delta) \\ &\leq 4 \left| \sum_{b \in \mathcal{N}} \mathrm{sign}(x,b) 2^{-|N(b)|} \right| \\ &\leq 4 \left| \sum_{b \in N_{\leq 1}(x,T[\ell-1])} \mathrm{sign}(x,b) 2^{-|N(b)|} \right| + 4 \sum_{x \in N_1(x,T[\ell-1])} 2^{-|N(b)|} \\ &\leq 10^{-6}, \end{aligned}$$

again for $\rho \geq \rho_0$ large. Combining the two above estimate, we obtain $|P_{\leq 1}\,(x \to a)\,[\ell] - 1| = 3 \cdot 10^{-6}$. Finally, as $\mu_{b\to x}(0)\,[\ell] > 0$ for all $b \in N(x) \setminus a$, we see that

$$\left| 1 - \frac{\mu_{x\to a}(1)\,[\ell]}{\mu_{x\to a}(0)\,[\ell]} \right| = |1 - P_{\leq 1}\,(x \to a)\,[\ell] \cdot P_{>1}\,(x \to a)\,[\ell]| \leq 10^{-4} + 4 \cdot 10^{-6} \leq 0.001.$$

Since $\mu_{x\to a}(1)\,[\ell] + \mu_{x\to a}(0)\,[\ell] = 1$ by construction, this entails the desired bound on $\Delta_{x\to a}\,[\ell]$. □

Corollary F.6 *Let b be a clause such that $N(b) \not\subset T\,[\ell]$. Let $x \in N(b)$. Then*

$$\mu_{b\to x}(0)\,[\ell-1] > 0 \textit{ and } \left| \frac{\mu_{b\to x}(1)\,[\ell-1]}{\mu_{b\to x}(0)\,[\ell-1]} - 1 \right| \leq \exp\left(-k_1/2\right).$$

Proof. We consider two cases.

Case 1: $|N(b) \setminus T\,[\ell-1]| > k_1$**.** Then the assertion is immediate from Lemma F.3.

Case 2: $|N(b) \setminus T\,[\ell-1]| \leq k_1$**.** Since $N(b) \not\subset T\,[\ell]$, condition **T2a** ensures that $0.1\theta k \leq |N(b)| \leq 10\theta k$. The assumption $N(b) \not\subset T\,[\ell]$ ensures that $b \notin T_3\,[\ell]$. Hence, the definition of $T_3\,[\ell]$ entails that $b \in T_3\,[\ell-1]$. As $N(b) \not\subset T_2\,[\ell]$, **T2e** implies

$$|N(b) \cap H\,[\ell-1]| \geq 3|N(b)|/4. \tag{38}$$

Let $J = N(b) \cap H\,[\ell-1]$. Since $b \in T_3\,[\ell-1]$, we have $\ell \geq 2$ and $|N(b) \setminus T\,[\ell-2]\,| \leq k_1$. Therefore, Corollary F.5 implies that

$$\left| \frac{\mu_{y\to b}(1)\,[\ell-1]}{\mu_{y\to b}(0)\,[\ell-1]} - 1 \right| \leq 0.1 \quad \text{for all } y \in J.$$

Consequently, (38) yields $\mu_{b\to x}(0)\,[\ell] > 0$ and

$$\left| \frac{\mu_{b\to x}(1)\,[\ell-1]}{\mu_{b\to x}(0)\,[\ell-1]} - 1 \right| \leq \left(\frac{2}{3}\right)^{|J|} \leq \left(\frac{2}{3}\right)^{|N(b)|/4} \leq \exp\left(-\theta k/10^5\right).$$

Thus, we have established the assertion in either case. □

Proof of Proposition F.1. We proceed by induction on ℓ. Since $B[0] = \emptyset$ the assertion is trivial for $\ell = 0$. Thus, assume that $\ell \geq 0$ and that $B[\ell] \subset T[\ell]$. Let $x \in V_t \setminus T[\ell+1]$. We will prove $x \notin B[\ell+1]$. Corollary F.6 implies that

$$\mu_{a\to x}(0)[\ell] > 0 \text{ and } \left|\frac{\mu_{a\to x}(1)[\ell]}{\mu_{a\to x}(0)[\ell]} - 1\right| \leq \exp(-k_1/2) \quad \text{for all } x \notin T[\ell+1],\, a \in N(x). \tag{39}$$

We claim

$$|P_{>1}(x \to a)[\ell+1] - 1| \leq \delta/100 \quad \text{for all } x \notin T[\ell+1],\, a \in N(x). \tag{40}$$

To establish (40), we consider two cases.

Case 1: $x \notin N(T_3[\ell])$**.** Let $\mathcal{T} = N_{>1}(x \to a)[\ell+1]$ be the set of all clauses b that contribute to the product $P_{>1}(x \to a)[\ell+1]$. Since $x \notin N(T_3[\ell] \cup T_3[\ell+1])$, none of the clauses $b \in \mathcal{T}$ features more than $|N(b)| - k_1$ variables from $T[\ell]$ (just by the definition of $T_3[\ell+1]$). Furthermore, because $x \notin T_2[\ell+1]$, **T2c** is not satisfied and thus we obtain the bound

$$\sum_{b\in\mathcal{T}} 2^{|N(b)\cap T[\ell]\setminus x| - |N(b)|} \leq \sum_{b\in N_{>1}(x,T[\ell])} 2^{|N(b)\cap T[\ell]\setminus x| - |N(b)|} \leq \delta/10^4, \tag{41}$$

provided that $\theta k \geq \rho \geq \rho_0$ is sufficiently large. Therefore, (40) follows from Corollary F.4.

Case 2: $x \in N(T_3[\ell])$**.** Let $\mathcal{T} = N_{>1}(x \to a)[\ell+1] \setminus T_3[\ell]$ be the set of all clauses b that occur in the product $P_{>1}(x \to a)[\ell+1]$, apart from the ones in $T_3[\ell]$. Since $x \notin T_2[\ell+1] \cup N(T_3[\ell+1])$, this set $\mathcal{T}$ also satisfies (41). Thus, Corollary F.4 yields

$$\left|1 - \prod_{b\in\mathcal{T}} \frac{\mu_{b\to x}(1)[\ell+1]}{\mu_{b\to x}(0)[\ell+1]}\right| \quad \leq \quad \delta/10^3. \tag{42}$$

As condition **T2d** in the definition of $T_2[\ell+1]$ ensures that $\left|N^1_{x\to a}[\ell+1] \cap T_3[\ell]\right| \leq 100$, (39) and (42) imply

$$\begin{aligned}
|1 - P_{>1}(x \to a)[\ell+1]| \quad &= \quad \left|1 - \prod_{b\in\mathcal{T}} \frac{\mu_{b\to x}(1)[\ell]}{\mu_{b\to x}(0)[\ell]} \cdot \prod_{b\in N_{>1}(x\to a)[\ell+1]\cap T_3[\ell]} \frac{\mu_{b\to x}(1)[\ell]}{\mu_{b\to x}(0)[\ell]}\right| \\
&= \quad \delta/10^{-3} + 2\left|1 - \prod_{b\in N^1_{x\to a}[\ell+1]\cap T_3[\ell]} \frac{\mu_{b\to x}(1)[\ell]}{\mu_{b\to x}(0)[\ell]}\right| \\
&\overset{(39)}{\leq} \quad \delta/10^{-3} + 2\left|1 - \prod_{b\in N^1_{x\to a}[\ell+1]\cap T_3[\ell]} 1 + \exp\left(-\theta k/10^5\right)\right| \\
&\leq \quad \delta/500.
\end{aligned}$$

which establishes (40).

If $x \notin T_1[\ell+1] \subset T[\ell+1]$ and $a \in N(x)$, then $|P_{\leq 1}(x \to a)[\ell+1] - 1| \leq \delta/100$. Hence, (40) implies that for all $x \notin T[\ell+1]$ and all $a \in N(x)$ we have $\mu_{x\to a}(0)[\ell+1] > 0$ and

$$\left|1 - \frac{\mu_{x\to a}(1)[\ell+1]}{\mu_{x\to a}(0)[\ell+1]}\right| = |1 - P_{\leq 1}(x \to a)[\ell+1] \cdot P_{>1}(x \to a)[\ell+1]| \leq \delta/99.$$

Consequently, $|\Delta_{x\to a}[\ell+1]| < 0.1\delta$, and thus $x \notin B[\ell+1]$. □

F.3 Proof of Proposition F.2

Lemma F.7 *Assume that* $|T_1\,[\ell] \cup T_2\,[\ell]| \le \theta\delta n/3$ *and* $|N(T_3\,[\ell])| \le \theta\delta n/2$. *Then* $|N(T_3\,[\ell+1])| \le \theta\delta n/2$.

Proof. By construction we have $T_3\,[\ell] \cap T_3\,[\ell+1] = \emptyset$. Furthermore, also by construction $N(T_3\,[\ell]) \subset T\,[\ell]$, and each clause in $T_3\,[\ell+1]$ has at least a 0.99-fraction of its variables in $T\,[\ell]$. Hence, **Q3** yields

$$\begin{aligned} |N(T_3\,[\ell+1])| + |N(T_3\,[\ell])| &\le \sum_{b \in T_3\,[\ell] \cup T_3\,[\ell+1]} |N(b)| \qquad [] \\ &\le 1.03|T\,[\ell]\,| \le 1.03\left(|T_1\,[\ell]| + |T_2\,[\ell]| + |N(T_3\,[\ell])|\right). \end{aligned}$$

Hence, $|N(T_3\,[\ell+1])| \le 1.03\,(|T_1\,[\ell]| + |T_2\,[\ell]|) + 0.03\,|N(T_3\,[\ell])| \le \theta\delta n/2$. □

Lemma F.8 *Assume that* $|T_1\,[\ell] \cup T_2\,[\ell]| \le \theta\delta n/3$ *and* $|N(T_3\,[\ell])| \le \theta\delta n/2$. *Moreover, suppose that* $|T\,[\ell-1]| \le \theta\delta n$. *Then* $|T_2\,[\ell+1]| \le \theta\delta n/6$.

Proof. Conditions **Q0** and **Q1** readily imply that the number of variables that satisfy either **T2a** or **T2b** is $\le 0.001\theta\delta n$. Moreover, we apply **Q2** to the set $T\,[\ell]$ of size

$$|T\,[\ell]| \le |T_1\,[\ell] \cup T_2\,[\ell]| + |N(T_3\,[\ell])| \le 0.9\theta\delta n \tag{43}$$

to conclude that the number of variables satisfying **T2c** is $\le 0.001\theta\delta n$ as well.

To bound the number of variables that satisfy **T2d**, consider the subgraph of the factor graph induced on $T_3\,[\ell] \cup N(T_3\,[\ell])$. For each $x \in N(T_3\,[\ell])$ let D_x be the number of neighbors of x in $T_3\,[\ell]$. Let ν be the set of all $x \in V_t$ so that $D_x \ge 2$. Then **Q3** yields

$$100\nu \quad \le \sum_{x \in N(T_3[\ell])} D_x = \sum_{a \in T_3[\ell]} |N(a)| \le 1.01|T\,[\ell]\,| \le \theta\delta n,$$

because $N(b) \subset T\,[\ell]$ for all $b \in T_3\,[\ell]$. Hence, there are at most $\nu \le 0.01\theta\delta n$ variables that satisfy **T2d**. In summary, we have shown that

$$|\{x \in V : x \text{ satisfies one of } \textbf{T2a–T2d}\}| \le 0.015\theta\delta n. \tag{44}$$

To deal with **T2e**, observe that if a clause a has at least $|N(a)|/4$ variables that are *not* harmless, then one of the following statements is true.

1. a contains at least $|N(a)|/20$ variables x that violate either **H1**, **H2**, or **H4**.
2. a contains at least $|N(a)|/5$ variables x that violate condition **H3**.

Let $\mathcal{C}_1$ be the set of clauses a for which 1. holds, and let $\mathcal{C}_2$ be the set of clauses satisfying 2., so that the number of variables satisfying **T2e** is bounded by $\sum_{a \in \mathcal{C}_1 \cup \mathcal{C}_2} |N(a)|$.

To bound $\mathcal{C}_1$, let Q be the set of all variables x that violate either **H1**, **H2**, or **H4** at time ℓ. Then conditions **Q1** and **Q2** entail that $|Q| \le 2 \cdot 10^{-4}\theta\delta n$ (because we are assuming $|T(\ell-1)| \le \theta\delta n$). Therefore, condition **Q3** implies that

$$\sum_{a \in \mathcal{C}_1} |N(a)| \le 21\,|Q| \le 0.003\theta\delta n. \tag{45}$$

To deal with $\mathcal{C}_2$ let $\mathcal{B}$ be the set of all clauses b such that $|N(b)| \ge 100k_1$ but $|N(b) \setminus T\,[\ell]\,| \le k_1$. Since we know from (43) that $|T\,[\ell]| \le \theta\delta n$, condition **Q3** applied to $T\,[\ell]$ implies that $\sum_{b \in \mathcal{B}} |N(b)| \le 1.03\,|T\,[\ell]|$. Thus,

$$|N(\mathcal{B})| \le \sum_{b \in \mathcal{B}} |N(b)| \le 1.03\,|T\,[\ell]| \le 1.03\theta\delta n. \tag{46}$$

Furthermore, let $\mathcal{U}$ be the set of all clauses a such that $N(a) \subset N(\mathcal{B})$. Let U be the set of variables $x \in N(\mathcal{B})$ that occur in at least two clauses from $\mathcal{U}$. Then by **Q3**

$$|U| + |N(\mathcal{B})| \quad \le \sum_{a \in \mathcal{U}} |N(a)| \le 1.01|N(\mathcal{B})|,$$

whence $|U| \leq 0.01|N(\mathcal{B})| \leq 0.02\theta\delta n$. Since $\mathcal{B} \subset \mathcal{U}$, the set U contains all variables that occur in at least two clauses from $\mathcal{B}$, i.e., all variables that violate condition **H3**. Therefore, any $a \in \mathcal{C}_2$ contains at least $|N(a)|/5$ variables from U. Applying **Q3** once more, we obtain

$$|N(\mathcal{C}_2)| \quad \leq \quad \sum_{a \in \mathcal{C}_2} |N(a)| \leq 5.05 \cdot 0.02\theta\delta n = 0.12\theta\delta n.$$

Combining this estimate with the bound (45) on $\mathcal{C}_1$, we conclude that the number of variables satisfying **T2e** is bounded by $\sum_{a \in \mathcal{C}_1 \cup \mathcal{C}_2} |N(a)| \leq 0.123\theta\delta n$. Together with (44) this yields the assertion. □

In Section F.4 we will derive the following bound on $|T_1(\ell+1)|$.

Proposition F.9 *If $|T[\ell]| \leq \theta\delta n$, then $|T_1(\ell+1) \setminus T_2(\ell+1)| \leq \theta\delta n/6$.*

Proof.[Proposition F.2] We are going to show that

$$|T_1[\ell] \cup T_2[\ell]| \leq \theta\delta n/3 \text{ and } |N(T_3[\ell])| \leq \theta\delta n/2 \tag{47}$$

for all $\ell \geq 0$. This implies that $|T[\ell]| \leq \theta\delta n$ for all $\ell \geq 0$, as desired.

In order to prove (47) we proceed by induction on ℓ. The bounds for $\ell = 0$ are direct from **Q0** and **Q1**. Now assume that (47) holds for all $l \leq \ell$. Then Lemma F.7 shows that $|N(T_3[\ell])| \leq \theta\delta n/2$. Moreover, Lemma F.8 applies, giving $|T_2[\ell+1]| \leq \theta\delta n/6$. Finally, Proposition F.9 shows $|T_1(\ell+1) \setminus T_2(\ell+1)| \leq \theta\delta n/6$, whence $|T_1[\ell+1] \cup T_2[\ell+1]| \leq \theta\delta n/3$. □

F.4 Proof of Proposition F.9

Throughout this section we assume that $|T[\ell]| \leq \theta\delta n$.

For a variable $x \in V_t$ and $a \in N(x)$ we let

$$\begin{aligned}
\sigma_{x\to a}[\ell+1] &= \sum_{b \in N_{\leq 1}(x\to a)[\ell+1]} 2^{1-|N(b)|}\mathrm{sign}(x,b), \\
\xi_{x\to a}[\ell+1] &= \sum_{b \in N_{\leq 1}(x\to a)[\ell+1]} \sum_{y \in N(b)\setminus x} 2^{1-|N(b)|}\mathrm{sign}(x,b)\,\mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell],
\end{aligned}$$

and $L_{x\to a}[\ell+1] = \sigma_{x\to a}[\ell+1] + \xi_{x\to a}[\ell+1]$. In Section F.5 we are going to establish the following.

Proposition F.10 *For any variable $x \notin T_2[\ell+1]$ and any clause $a \in N(x)$ we have*

$$\left|L_{x\to a}[\ell+1] + \ln P^0_{x\to a}[\ell+1]\right| \leq 10^{-3}\delta.$$

Lemma F.11 *For all but at most $0.01\theta\delta n$ variables $x \in V \setminus T_2[\ell+1]$ we have*

$$\max_{a \in N(x)} |\sigma_{x\to a}[\ell+1]| \leq 0.001\delta.$$

Proof. This is a direct consequence of condition **Q2**. □

Lemma F.12 *Let x be a variable and let $b_1, b_2 \in N(x)$ be such that $|N(b_i) \cap T[\ell]| \leq 2$ and $|N(b_i)| \geq 0.1\theta k$ for $i = 1, 2$. Then $|\Delta_{x\to b_1}[\ell] - \Delta_{x\to b_2}[\ell]| \leq \delta^3$.*

Proof. Our assumptions ensure that $N(b_i) \setminus T \neq \emptyset$ and thus $\mu_{x\to b_i}(0)[\ell] > 0$ for $i = 1, 2$. Corollary F.6 yields

$$\mu_{b_i\to x}(0)[\ell-1] > 0 \text{ and } \left|\frac{\mu_{b_i\to x}(1)[\ell-1]}{\mu_{b_i\to x}(0)[\ell-1]} - 1\right| \leq \exp(-k_1/2) \tag{48}$$

for $i = 1, 2$. Let $q_i = \frac{\mu_{x\to b_i}(1)[\ell]}{\mu_{x\to b_i}(0)[\ell]}$. Then

$$\begin{aligned}
|\Delta_{x\to b_1}[\ell] - \Delta_{x\to b_2}[\ell]| &= |\mu_{x\to b_1}(1)[\ell] - \mu_{x\to b_2}(1)[\ell]| \\
&= \left|\frac{q_1 - q_2}{(1+q_1)(1+q_2)}\right| \leq \left|1 - \frac{q_2}{q_1}\right| \\
&= \left|1 - \frac{\mu_{b_1\to x}(1)[\ell-1]}{\mu_{b_1\to x}(0)[\ell-1]} \cdot \frac{\mu_{b_2\to x}(0)[\ell-1]}{\mu_{b_2\to x}(1)[\ell-1]}\right|.
\end{aligned} \tag{49}$$

Hence, the assertion follows from (48), provided that the constant c in the definition (6) of $\delta = \delta_t$ is chosen sufficiently small. □

Lemma F.13 *For all but at most $0.1\theta\delta n$ variables $x \notin T_2[\ell+1]$ we have $\max_{a \in N(x)} |\xi_{x\to a}| \le 0.001\delta$.*

Proof. For a variable y let $\mathcal{N}(y)$ be the set of all clauses $b \in N(y)$ such that $b \in N_{\le 1}(x, T[\ell])$ for some variable $x \in V$. If $\mathcal{N}(y) = \emptyset$ we define $\Delta_y = 0$; otherwise select $a_y \in \mathcal{N}(y)$ arbitrarily and set $\Delta_y = \Delta_{y \to a_y}[\ell]$. Thus, we obtain a vector $\Delta = (\Delta_y)_{y\in V}$ with norm $\|\Delta\|_\infty \le \frac{1}{2}$. Let $\Xi = (\xi_x)_{x\in V} = \Lambda_{T[\ell]}\Delta$, where $\Lambda_{T[\ell]}$ is the linear operator from condition **Q4** in Definition 4.1. That is, for any $x \in V$ we have

$$\xi_x = \sum_{b \in N_{\le 1}(x, T[\ell])} \sum_{y \in N(b)\setminus x} 2^{1-|N(b)|}\mathrm{sign}(x,b)\mathrm{sign}(y,b)\Delta_y.$$

Because $|T[\ell]| \le \theta\delta n$, condition **Q4** ensures that $\left\|\Lambda_{T[\ell]}\right\|_\square \le \delta^4\theta n$. Consequently,

$$\|\Xi\|_1 = \left\|\Lambda_{T[\ell]}\Delta\right\|_1 \le \left\|\Lambda_{T[\ell]}\right\|_\square \|\Delta\|_\infty \le \delta^4\theta n. \tag{50}$$

Since $\|\Xi\|_1 = \sum_{x\in V}|\xi_x|$, (50) implies that

$$\left|\left\{x \in V : |\xi_x| > \delta^2\right\}\right| \le \delta^2\theta n. \tag{51}$$

To infer the lemma from (51), we need to establish a relation between ξ_x and $\xi_{x\to a}$ for $x \in V \setminus T_2[\ell+1]$ and $a \in N(x)$. Applying Lemma F.12, we get

$$|\Delta_{y\to b}[\ell] - \Delta_{y\to b'}[\ell]| \le \delta^3 \quad \text{for all } y \in V,\ b, b' \in \mathcal{N}(y). \tag{52}$$

Hence,

$$|\Delta_{y\to b}[\ell] - \Delta_y| \le \delta^2. \tag{53}$$

Consequently, we obtain for all $x \notin T_2[\ell+1]$

$$\begin{aligned}
\max_{a\in N(x)} |\xi_{x\to a} - \xi_x| &\le \sum_{y \in N(a)} 2^{1-|N(a)|}|\Delta_y| + \sum_{b \in N_{\le 1}(x\to a)[\ell+1]} \sum_{y\in N(b)\setminus x} 2^{1-|N(b)|}|\Delta_{y\to a}[\ell] - \Delta_y| \\
&\overset{(53)}{\le} |N(a)|2^{-|N(a)|} + 2\delta^3 \sum_{b\in N_{\le 1}(x\to a)[\ell+1]} |N(b)|2^{1-|N(b)|} \\
&\le 10^{-4}\delta \qquad \text{[by } \mathbf{Q2}\text{]},
\end{aligned} \tag{54}$$

provided that $\theta k \ge \rho \ge \rho_0$ is sufficiently large. If $x \in V \setminus T_2[\ell+1]$ is such that $|\xi_x| \le \delta^2$, then (54) implies that $|\xi_{x\to a}| \le 2\cdot 10^{-4}\delta$ for any $a \in N(x)$. Therefore, the assertion follows from (51). □

Proof.[Proposition F.9] Let S be the set of all variables $x \notin T_2[\ell+1]$ such that $\max_{a\in N(x)}|\sigma_{x\to a}| \le 0.001\delta$ and $\max_{a\in N(i)}|\xi_{x\to a}| \le 0.001\delta$. Then Proposition F.10 entails that for any $x \in S$ and $a \in N(x)$

$$|\ln P_{\le 1}(x\to a)[\ell+1]| \le |L_{x\to a}[\ell+1]| + 10^{-3}\delta = \sigma_{x\to a} + \xi_{x\to a} + 10^{-3}\delta \le 0.003\delta.$$

Hence, $|P_{\le 1}(x \to a) - 1| \le 0.01\delta$ for all $x \in S$, $a \in N(x)$, and therefore

$$T_1[\ell+1] \setminus T_2[\ell+1] \subset V \setminus (S \cup T_2[\ell+1])).$$

Finally, Lemmas F.11 and F.13 imply $|T_1[\ell+1] \setminus T_2[\ell+1]| \le |V \setminus (S \cup T_2[\ell+1])| \le \theta\delta n/6$. □

F.5 Proof of Proposition F.10

Lemma F.14 *Let x be a variable, $a \in N(x)$, and $b \in N_{\le 1}(x\to a)[\ell+1]$. Then $\mu_{b\to x}(0)[\ell] > 0$ and*

$$\begin{aligned}
\ln\left(\frac{\mu_{b\to x}(1)[\ell]}{\mu_{b\to x}(0)[\ell]}\right) &= -2^{1-|N(b)|}\left[\mathrm{sign}(x,b) + 2\sum_{y\in N(b)\setminus x}\mathrm{sign}(x,b)\,\mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell]\right] \\
&\quad + \frac{(\theta k\delta + |N(b)\cap T[\ell]\setminus x|)}{2^{|N(b)|}}\cdot O(\theta k\delta),
\end{aligned} \tag{55}$$

where the $O(\cdot)$ denotes an asymptotic bound for ρ sufficiently large.

Proof. If $b \in N_{\leq 1}(x \to a)[\ell+1]$, then $|N(b) \cap T[\ell]| \leq 2$, while $|N(b)| \geq 0.1\theta k$. Therefore, Lemma F.3 shows that $|\frac{1}{2} - \mu_{b\to x}(0)[\ell]| \leq \delta^2$, provided that the constant c in the definition (6) of δ is chosen sufficiently small. Furthermore, b is not redundant, and thus not a tautology, because otherwise $N(b) \subset T_2[\ell]$ due to **T2a**. In order to establish (55) we consider two cases. The first case is that $\mathrm{sign}(x,b) = -1$. Then $\mu_{b\to x}(0)[\ell] = 1$ and thus by definition (2)

$$\begin{aligned}
\frac{\mu_{b\to x}(1)[\ell]}{\mu_{b\to x}(0)[\ell]} &= \mu_{b\to x}(1)[\ell] = 1 - \prod_{y\in N(b)\setminus x} \mu_{y\to b}((1-\mathrm{sign}(y,b))/2)[\ell] \\
&= 1 - \prod_{y\in N(b)\setminus x} \frac{1}{2} - \mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] \\
&= 1 - 2^{1-|N(b)|} \prod_{y\in N(b)\setminus x} 1 - 2\mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell]. \qquad (56)
\end{aligned}$$

Let $\Gamma = N(b) \setminus (T[\ell] \cup x)$. As Proposition F.1 shows that $T[\ell] \supset B[\ell]$ contains all biased variables, we have $|\Delta_{y\to b}[\ell]| \leq \delta$ for all $y \in \Gamma$. Therefore, we can use the approximation $|\ln(1-z) + z| \leq z^2$ for $|z| \leq \frac{1}{2}$ to obtain

$$\begin{aligned}
&\left| \left[\ln \prod_{y\in\Gamma} 1 - 2\mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] \right] + \sum_{y\in\Gamma} 2\mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] \right| \\
&\qquad \leq 4 \sum_{y\in\Gamma} \Delta_{y\to b}[\ell]^2 \leq 40\theta k \delta^2; \qquad (57)
\end{aligned}$$

in the last step, we used that $|N(b)| \leq 10\theta k$ for all $b \in N_{\leq 1}(x \to a)[\ell+1]$. Furthermore, by the definition of $N_{\leq 1}(x \to a)$, we have $|N(b) \cap T[\ell] \setminus x| \leq 1$. If there is $y_0 \in N(b) \cap T[\ell] \setminus x$, then (57) yields

$$\begin{aligned}
\prod_{y\in N(b)\setminus x} 1 - 2\mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] &= (1 - 2\mathrm{sign}(y_0,b)\,\Delta_{y_0\to b}[\ell]) \cdot \prod_{y\in\Gamma} 1 - 2\mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] \\
&= 1 - 2\left[\sum_{y\in N(b)\setminus x} \mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] \right] + O(\theta k \delta)^2 \\
&\qquad -4\mathrm{sign}(y_0,b)\,\Delta_{y_0\to b}[\ell] \sum_{y\in N(b)\setminus x} \mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] \\
&= 1 - 2\left[\sum_{y\in N(b)\setminus x} \mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] \right] + O(\theta k \delta),
\end{aligned}$$

where the $O(\cdot)$ hides a term that remains bounded independently of ρ. Hence, in any case we have

$$\begin{aligned}
\prod_{y\in N(b)\setminus x} 1 - 2\mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] &= 1 - 2\left[\sum_{y\in N(b)\setminus x} \mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] \right] \\
&\qquad +(\theta k \delta + |N(b) \cap T[\ell] \setminus x|) \cdot O(\theta k \delta).
\end{aligned}$$

Combining this with (56) and using the approximation $|\ln(1-z) + z| \leq z^2$ for $|z| \leq 1/2$, we see that

$$\begin{aligned}
\ln\left(\frac{\mu_{b\to x}(1)[\ell]}{\mu_{b\to x}(0)[\ell]} \right) &= -2^{1-|N(b)|} \left[\prod_{y\in N(b)\setminus x} 1 - 2\mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] \right] + O(2^{2-2|N(b)|}) \\
&= -2^{1-|N(b)|} \left[\mathrm{sign}(x,b) + 2 \sum_{y\in N(b)\setminus x} \mathrm{sign}(x,b)\,\mathrm{sign}(y,b)\,\Delta_{y\to b}[\ell] \right] \\
&\qquad +2^{1-|N(b)|}(\theta k \delta + |N(b) \cap T[\ell] \setminus x|) \cdot O(\theta k \delta).
\end{aligned}$$

In the case $\mathrm{sign}(x,a) = 1$ the proof is similar. □

Proof.[Proposition F.10] Lemma F.14 entails that for large ρ

$$\begin{aligned}
\ln P_{\leq 1}\left(x\to a\right)\left[\ell+1\right] &= \sum_{b\in N_{\leq 1}(x\to a)[\ell+1]} \ln\left(\frac{\mu_{b\to x}(1)\left[\ell\right]}{\mu_{b\to x}(0)\left[\ell\right]}\right)\\
&= \sum_{b\in N_{\leq 1}(x\to a)[\ell+1]} -L_{b\to x}\left[\ell+1\right]+2^{-|N(b)|}(\theta k\delta+|N(b)\cap T\left[\ell\right]\setminus x|)O(\theta k\delta)\\
&= \sum_{b\in N_{\leq 1}(x\to a)[\ell+1]} -L_{b\to x}\left[\ell+1\right]+O(\theta k\delta)^2\sum_{b\in N^0_{x\to a}[\ell+1]} 2^{-|N(b)|}\\
&\qquad +O(\theta k\delta)\sum_{b\in N_1(x\to a)[\ell+1]\cap N_{\leq 1}(x\to a)[\ell+1]} 2^{-|N(b)|}\\
&= -\sum_{b\in N_{\leq 1}(x\to a)[\ell+1]} L_{b\to x}\left[\ell+1\right]+\frac{O(\theta k\delta)^2}{(\theta k)^3\delta}+O(\theta k\delta)\cdot O(\rho\delta)\\
&= -\sum_{b\in N_{\leq 1}(x\to a)[\ell+1]} L_{b\to x}\left[\ell+1\right]+o(\delta),
\end{aligned}$$

due to conditions **T2b** and **T2c**. Hence, there is $\rho_0>0$ such that for $\rho>\rho_0$ we have

$$\left|\sum_{b\in N_{\leq 1}(x\to a)[\ell+1]} L_{b\to x}\left[\ell+1\right]+\ln P_{\leq 1}\left(x\to a\right)\left[\ell+1\right]\right| \quad\leq\quad 10^{-3}\delta,$$

as desired. $\square$

F.6 Completing the proof of Theorem 4.3

We are going to show that $\left|\mu_x(\Phi_{t-1},\omega)-\frac{1}{2}\right|\leq\delta=\delta_t$ for all $x\in V_t\setminus T\left[\omega+1\right]$. This will imply Theorem 4.3, because $|T\left[\omega+1\right]|\leq\delta_t(n-t)$ by Proposition F.2.

Thus, let $x\in V_t\setminus T\left[\omega+1\right]$. Corollary F.6 shows that $\mu_{b\to x}(\zeta)\left[\omega\right]>0$ for $\zeta=0,1$. Hence,

$$P(1)=\prod_{b\in N(x)}\mu_{b\to x}(1)\left[\omega\right]>0,\quad\text{and}\quad P(0)=\prod_{b\in N(x)}\mu_{b\to x}(0)\left[\omega\right]>0.$$

Recall from (4) that

$$\mu_x(\Phi_{t-1},\omega)=\frac{P(1)}{P(0)+P(1)}.$$

If $N(x)=\emptyset$, then trivially $P(0)=P(1)=1$ and thus $\mu_x(\Phi_{t-1},\omega)=\frac{1}{2}$. Thus, assume that $N(x)\neq\emptyset$ and pick an arbitrary $a\in N(x)$. Then by (3)

$$P(\zeta)\quad=\quad\mu_{a\to x}(\zeta)\left[\omega\right]\cdot\mu_{x\to a}(\zeta)\left[\omega+1\right]\qquad\text{for }\zeta=0,1.$$

Since $x\notin T\left[\omega+1\right]\supset B\left[\omega+1\right]$ (by Proposition F.1), we have

$$\left|\mu_{x\to a}(\zeta)\left[\omega+1\right]-\frac{1}{2}\right|\leq 0.1\delta$$

for $\zeta=0,1$. Therefore, assuming that $\theta k\geq\rho\geq\rho_0$ is sufficiently large and thus $\delta=\exp(-c\theta k)$ is sufficiently small, we obtain

$$\begin{aligned}
\ln\frac{\mu_{x\to a}(0)\left[\omega+1\right]}{\mu_{x\to a}(1)\left[\omega+1\right]} &\leq \ln\frac{1+0.2\delta}{1-0.2\delta}\leq 0.5\delta,\qquad\text{and analogously}\\
\ln\frac{\mu_{x\to a}(0)\left[\omega+1\right]}{\mu_{x\to a}(1)\left[\omega+1\right]} &\geq -0.5\delta.
\end{aligned}$$

Furthermore, Corollary F.6 yields

$$\left|\ln\frac{\mu_{a\to x}(0)\left[\omega\right]}{\mu_{a\to x}(1)\left[\omega\right]}\right|\leq 2\exp(-\theta k/10^5)\leq\delta^2.$$

Hence,

$$\left|\ln\frac{P(0)}{P(1)}\right| \quad\leq\quad \left|\ln\frac{\mu_{x\to a}(0)\,[\omega+1]}{\mu_{x\to a}(1)\,[\omega+1]}\right| + \left|\ln\frac{\mu_{a\to x}(0)\,[\omega]}{\mu_{a\to x}(1)\,[\omega]}\right| \leq 0.5\delta+\delta^2 \leq 0.51\delta.$$

Therefore, letting $z=\ln\frac{P(0)}{P(1)}$, we obtain

$$\left|\frac{1}{2}-\mu_x(\Phi_{t-1},\omega)\right| \quad=\quad \left|\frac{1}{2}-\frac{P(1)}{P(0)+P(1)}\right| = \left|\frac{1}{2}-\frac{1}{2}\cdot\frac{1}{1+\exp(z)/2}\right| \leq \exp(z) \leq 0.51\delta,$$

as desired.

G Proof of Corollary 1.6

Let us begin by disproving the second statement of Hypothesis 1.3. Suppose that r is as in (1). Let us call a formula Φ (ω,ε)-***simple*** if

$$\mathrm{P}\left[\forall t\leq n: |M_{x_t}(\Phi_t)-\mu_{x_t}(\Phi_t,\omega)|\leq\varepsilon\right]\geq 1-o(1), \tag{58}$$

where, of course, probability is taken over the execution of Experiment 1.2 only. Then Hypothesis 1.3 claims that for any $\varepsilon>0$ there exist a function $\omega=\omega_\varepsilon(n)$ and a number $\eta=\eta(\varepsilon,k,r)$ such that the random formula $\mathbf{\Phi}$ is (ω,ε)-simple with probability at least η. (Actually, Hypothesis 1.3 claims that this statement is true for *any* fixed $\eta<1$.) Assume for contradiction that this statement is true.

Let σ_{ideal} be the (random) assignment generated by Experiment 1.2, and let σ_{BP} be the assignment constructed by `BPdec` on the same random input $\mathbf{\Phi}$. Furthermore, let $\mathcal{G}=\mathcal{G}(\Phi)$ be the set of all satisfying assignments σ such that given $\sigma_{\mathrm{ideal}}=\sigma$ it is true that $|M_{x_t}(\Phi_{t-1})-\mu_{x_t}(\Phi_{t-1},\omega)|\leq\varepsilon$ w.h.p. for all $1\leq t\leq n$. Then, by assumption (58), for a random formula $\mathbf{\Phi}$ we have $|\mathcal{G}|=(1-o(1))2^n$ with probability at least $\eta>0$.

Let Φ be a formula such that indeed $|\mathcal{G}(\Phi)|=(1-o(1))2^n$ w.h.p., and let $\sigma\in\mathcal{G}(\Phi)$. Let $\mathcal{E}_t$ be the event that there is an $\sigma\in\mathcal{G}$ such that $\sigma_{\mathrm{ideal}}(x_s)=\sigma_{\mathrm{BP}}(x_s)=\sigma(x_s)$ for all $1\leq s\leq t$. Given that $\mathcal{E}_{t-1}$ occurs, we have $|M_{x_t}(\Phi_{t-1})-\mu_{x_t}(\Phi_{t-1},\omega)|\leq\varepsilon$. Hence, by Steps 3 of Experiment 1.2 and `BPdec` and because $|\mathcal{G}|=(1-o(1))2^n$ we have $\mathrm{P}\left[\mathcal{E}_t|\mathcal{E}_{t-1}\right]\geq 1-\varepsilon$. Multiplying these terms up, we see that

$$\mathrm{P}\left[\mathcal{E}_n\right]\geq\exp(-\varepsilon n). \tag{59}$$

But if $\mathcal{E}_n$ occurs, then the output σ_{BP} is a satisfying assignment. Hence, Theorem 1.1 implies that there is an $\varepsilon_0=\varepsilon_0(k,r)>0$ (namely, the constant hidden in the $\Omega(\cdot)$ in Theorem 1.1) such that (59) is false for $\varepsilon<\varepsilon_0$, in contradiction to part ii. of Hypothesis 1.3.

We come to the first statement. Assume for contradiction that there is a function $\gamma=\gamma(n)\to\infty$ so that with probability $1-\exp(-\gamma)$ the random formula $\mathbf{\Phi}$ is such that for the sequence $\mathbf{\Phi}_t$ generated by Experiment 1.2 we have

$$\mathrm{P}\left[\forall t\leq n: |M_{x_t}(\mathbf{\Phi}_t)-M_{x_t}(\mathbf{\Phi}_t,\omega)|\leq\varepsilon\right]\geq 1-\exp(-\gamma) \tag{60}$$

for any $\varepsilon>0$ and $\omega=\omega(\varepsilon,k,r)>1$ sufficiently large.

Lemma G.1 *There is a number $\eta=\eta(\varepsilon,k,r)>0$ such that with probability $\geq\eta$ the factor graph of $\mathbf{\Phi}$ has no cycle of length $\leq 2\omega$.*

Lemma G.1 is a standard result in the study of random hypergraphs. It follows easily from the fact that the total number of cycles of length $\leq 2\omega$ converges in distribution to a Poisson distribution with mean $\exp(C\omega)$ for a certain $C=C(k,r)>0$ [17, Section 9.2]. Therefore, the probability that there is no such cycle converges to $\eta(\varepsilon,k,r)=\exp(-\exp(C\omega))>0$.

Let $\mathcal{A}$ be the event that the factor graph of $\mathbf{\Phi}$ has no cycle of length $\leq\omega$. Then (60) entails that for large enough n the formula $\mathbf{\Phi}$ satisfies (60) with probability at least $1-\exp(-\gamma)/\eta=1-o(1)$ even if we condition on the event $\mathcal{A}$. Furthermore, if the event $\mathcal{A}$ occurs, then $M_{x_t}(\mathbf{\Phi}_t,\omega)=\mu_{x_t}(\mathbf{\Phi}_{t-1},\omega)$ because Belief Propagation yields the correct marginals if there is no cycle of length $\leq 2\omega$. Thus, we conclude that the with probability at least $\eta-o(1)$ the random formula $\mathbf{\Phi}$ is such that

$$\mathrm{P}\left[\forall t\leq n: |M_{x_t}(\mathbf{\Phi}_t)-\mu_{x_t}(\mathbf{\Phi}_t,\omega)|\leq\varepsilon\right]\geq 1-o(1), \tag{61}$$

which is precisely the statement that we refuted in the first part of the proof.